

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

OPTi Inc.

Plaintiff,

v.

NVIDIA Corporation

Defendant.

§
§
§
§
§
§
§
§
§
§

CIVIL ACTION 2:04-CV-377 (TJW)

JURY

**CONTAINS CONFIDENTIAL AND "OUTSIDE COUNSEL EYES ONLY"
INFORMATION OF THIRD PARTIES (INTEL, COMPAQ/HP, AND FREESCALE)**

DESIGNATED UNDER THE PROTECTIVE ORDER ENTERED IN THIS CASE

TABLE OF CONTENTS

	PAGE
I. INTRODUCTION	1
II. QUALIFICATIONS	1
III. PRIOR TESTIMONY AND COMPENSATION.....	2
IV. MATERIALS REVIEWED	3
V. ANALYSIS OF COURT’S CLAIM CONSTRUCTION RULING.....	5
VI. BACKGROUND INFORMATION	8
VII. PERSON OF ORDINARY SKILL IN THE ART	9
VIII. INSTRUCTIONS RE: LEGAL STANDARDS	10
A. Burden of proof instruction.....	10
B. Anticipation instructions	11
C. Prior invention instructions.....	11
D. Obviousness instructions	13
IX. SUMMARY OF OPINIONS	16
A. The asserted claims of the OPTi '906 presnoop patent are invalid in view of the prior invention of Intel employees, including Bruce Young, Ali Oztaskin, and Brian Langendorf (and also published in Shanley).....	16
B. The asserted claims of the OPTi '906 patent are invalid in view of the prior invention of Compaq employees, including Michael Collins and Gary Thome	18
C. The asserted claims of the OPTi '906 patent are invalid in view of the prior invention of Motorola employees, including Michael Garcia, Christopher Bryant, and Brian Reynolds.....	19
D. The asserted claims of the OPTi presnoop patents are invalid in view of the prior invention of IBM (Microchannel and Amini)	19
X. DETAILED DESCRIPTION OF AND BASIS FOR OPINIONS	19
A. The asserted claims of the OPTi patents are invalid in view of the prior invention by Mr. Young, Mr. Oztaskin, Mr. Hayek, and Mr. Langendorf (including the Shanley publication)	19

TABLE OF CONTENTS
(CONTINUED)

TABLE OF CONTENTS
(CONTINUED)

PAGE

B.	The asserted claims of the OPTi presnoop patents are invalid in view of the prior invention of Compaq employees, including Michael Collins and Gary Thome	141
1.	Compaq was a member of the PCI special interest group (PCI SIG)....	145
2.	The CMC/DMC architecture included a read prefetch buffer that could hold up to two cache lines in the bridge, allowing the CMC/DMC to burst across cache line boundaries without waiting for snoops at cache line boundaries	145
3.	The CMC/DMC, as described in the Thome/Collins patent, performed a “pipelining” process that both presnooped and prefetched.....	149
4.	The CMC/DMC, as reduced to practice in silicon and as described in the Thome/Collins patent, performed a “pipelining” process that both presnooped and prefetched	154
5.	The CMC/DMC chipset, as described in the Thome/Collins patent and as reduced to practice in silicon, did not disconnect at cache line boundaries, but limited bursts only to entire pages of memory and/or 256 addresses	169

TABLE OF CONTENTS
(CONTINUED)

	PAGE
6. The evidence supports Mr. Thome's and Mr. Collins' testimony that they had working silicon, able to practice the MRM command with presnooping and prefetching, in February 1994	173
7. Mr. Thome and Mr. Collins conceived of the presnooping and prefetching invention in the 1992-93 time frame, before OPTi.....	185
8. The documents that OPTi cites do not contradict Compaq's conception	190
9. The documents that OPTi cites do not contradict Compaq's reduction to practice.....	195
10. Compaq conceived and reduced to practice over a lengthy period of time because it was making a very complex design, much more sophisticated than the relatively simple presnoop-only method described in the OPTi patent.....	198
11. The Kramer Report shows that the Prosignia operated as the Compaq/HP witnesses testified.....	199
12. Acacia's analysis of the Compaq prior art is consistent with the rest of the evidence	202
13. Compaq claim chart & timeline	203
C. The asserted claims of the OPTi patents are invalid in view of the prior invention by Motorola inventors of the Motorola MPC105, and the MPC105 articles.....	203
1. The design of the MPC105 began in late 1992 or early 1993	204
2. The design of the MPC105 had been completed by at least March of 1994 when Motorola began submitting papers describing the MPC105 to academic and commercial conferences	205
3. Motorola received a working silicon chip for the MPC105 on May 22, 1994.....	210
4. After reducing the MPC105 to silicon, Motorola continued to publish articles and prepared to sell the MPC105 commercially	211
5. Motorola presented the MPC105 at Hot Chips VI in mid-August 1994 and formerly announced the MPC105 on August 22, 1994.....	212
6. Motorola made additional presentations of the MPC105 in the latter half of 1994 and expected to begin offering the MPC105 for sale at the end of 1994	215
7. The Motorola MPC105 was designed to operate as an interface between PowerPC 60x processors and a PCI bus	217

TABLE OF CONTENTS
(CONTINUED)

	PAGE
8. The MPC105 relied upon and implemented portions of the PCI Local Bus Specification, Revision 2.0, and also implemented presnooping as described in the Patents-in-Suit	219
9. The MPC105 had only one prefetch buffer, so it had to disconnect at cache line boundaries	223
10. The MPC105's disconnect at cache line boundaries is not relevant to several of the '906 claims, and even under OPTi's interpretation, '906 claims 9 and 26 would be obvious in view of the MPC105 system	228
11. The MPC105, as described by the evidence, was faster than the Viper described in the OPTi '906 patent	228
12. Summary of opinions to be expressed with respect to the Motorola MPC105	230
13. Claim chart & timeline.....	230
D. The IBM prior art (MicroChannel TDB and Amini).....	230
1. IBM Microchannel TDB (1993)	230
2. Amini (1993).....	231
E. In addition to being anticipated, the claims in suit were obvious in view of the Intel (including Shanley), Compaq, Motorola and IBM prior art.....	232
1. Differences between the invention and the prior art.....	232
2. Level of ordinary skill.....	233
3. The objective indications of obviousness support my opinion	233
a. The OPTi Viper was not a commercial success—Intel's Triton, which included both prefetching and presnooping, had far greater performance and was successful.....	234
b. No long-felt need existed – the PCI specification was very new and bus masters did not initially use the MRM command.....	235
c. Others did not try and fail to make the invention	236
d. No evidence suggests that anyone copied the Viper – OPTi trailed Intel and others in PCI chipset performance.....	236
e. The invention did not achieve unexpected results	237
f. The invention did not receive praise of others.....	237
g. The taking of licenses under the patent by others.....	238

TABLE OF CONTENTS
(CONTINUED)

	PAGE
h. No experts expressed surprise at the making of the invention—presnoop was fully expected and inevitable	240
i. The patentee did not proceed contrary to the accepted wisdom of the prior art.....	240
F. OPTi’s embodiment of the invention was a generation behind Intel and Compaq	241
G. Any claim that covers snooping ahead more than one line is not supported by a written description in the ‘906 patent specification	248
1. The Transfer of Line N Cannot Occur Until the Snoop of Line N is Completed	248
2. The Snoop of Line N+1 Cannot Occur Until the Snoop of Line N is Completed	249
3. The Transfer of Line N+1 Cannot Occur on the Bus Until the Transfer of Line N is Completed	249
4. OPTi’s Patent Teaches Away from Snooping Ahead More than One Line.....	250
XI. CONCLUSION AND RESERVATION	255

I. INTRODUCTION

1. My name is Robert P. Colwell. My address is 3594 N.W. Bronson Crest Loop, Portland, Oregon, 97229. I am over eighteen years of age and would otherwise be competent to testify to the matters in this report, if I am called upon to do so at trial.

2. I have been retained by Cooley Godward LLP on behalf of Defendant NVIDIA Corporation ("NVIDIA") to serve as an expert in the above-referenced matter.

3. For purposes of this report, I have been asked to provide an expert technical analysis and opinion as to whether the asserted claims of U.S. Patent No. 5,710,906 ("the '906 patent") are invalid. My opinions are based upon the information I have considered, received, and reviewed to date. My opinions may be supplemented or modified if I receive additional information, and may also be supplemented to reply to additional information that arises at trial or that is provided by OPTi, Inc. ("OPTi") or witnesses retained by OPTi.

II. QUALIFICATIONS

4. I received a Ph.D. in Computer Engineering in 1985, and an M.S.E.E. in Computer Engineering in 1978, from Carnegie-Mellon University. I received a B.S. in Electrical Engineering from the University of Pittsburgh in 1977.

5. From 1977-1980, I worked as a hardware engineer at Bell Telephone Labs in Holmdel, New Jersey working on 8-bit and 32-bit microprocessors. From 1981-1984, I worked as a hardware engineer for Perq Systems in Pittsburgh, Pennsylvania designing graphics-display hardware for first-generation engineering workstations. From 1985-1990, I worked as a hardware architect for Multiflow Computer in New Haven, Connecticut, where I was one of the seven hardware designers who created the world's first VLIW (very long instruction word) scientific supercomputer.

6. In 1990, Intel Corporation hired me as a senior CPU architect charged with designing Intel microprocessors. From 1990-1992, I was one of the three senior architects responsible for designing Intel's P6 microprocessor architecture, which is the processing core of

the well-known Pentium II, Pentium III, Celeron, Xeon, and Centrino families of Intel microprocessors. These microprocessors are commonly found in laptop, desktop, and server computer systems.

7. From 1992-2001, I was the lead Intel IA32 architect responsible for all of Intel's CPU architecture development efforts. My responsibilities included initiating and leading the development of Intel's Pentium 4 CPU core development, during which time I managed up to 100 other Intel engineers and employees. I was named an Intel Fellow in 1997. At that time, I was one of only 27 Intel Fellows out of approximately 80,000 Intel employees.

8. Since leaving Intel in 2001, I have been a consultant to industry and academia on various computer hardware and software issues. I have also written numerous publications, magazine articles, and authoritative texts in the field of CPU design and computer architecture. These publications include the P6 chapter of "Superscalar Processor Design" by Shen & Lipasti, published by McGraw-Hill in 2003, and numerous articles in IEEE Computer Magazine and other technical publications. I also frequently present talks at various symposia regarding CPU architecture and design, as well as computer systems in general.

9. I am a named inventor on numerous patents, including 35 patents that I worked on while at Intel. My work experience, patents, and awards are covered in more detail in my curriculum vitae, which is attached as Attachment A.

III. PRIOR TESTIMONY AND COMPENSATION

10. A list of my prior deposition and trial testimony over the last four years is attached as Attachment A. I am being compensated for my time on this case at a rate of \$250 per hour.

IV. MATERIALS REVIEWED

11. The materials that I have considered in connection with this report include, in addition to all material cited and referenced in the body of my report and in my charts and timelines:

- The OPTi presnoop patents '906, '036, and '291 patents. (EX. 17, 18, 19).
- The Prosecution Histories of the OPTi presnoop patents. (EX. 20, 21, 22).
- PCI Local Bus Specification Revision 2.0 (April 30, 1993) (EX. 64).
- Intel, CDC Cache & DRAM Controller Target Specification, Rev. 1.0 (EX. 122).
- Intel CDC Cache & DRAM Controller Target Specification, Rev. 0.2 (EX. 123).
- The Intel '094 Patent (EX. 96).
- Triton Host Memory/Cache Controller/PCI Bridge, Preliminary Component Specification Draft (EX. 98).
- The Triton System Controller (TSC) and Triton Data Path (TDP) Component Specification (EX. 99).
- Intel Invention Disclosure of Hayek, Oztaskin, Young, and Langendorf (EX. 100).
- Report of Barry L. Kramer re: Intel Triton analysis (Attachment B) (I also conferred with Mr. Kramer about his analysis).
- Report of Barry L. Kramer re: Compaq Prosignia 500 analysis (Attachment C) (I also conferred with Mr. Kramer about his analysis).
- Report of Gerald Mossinghoff.
- Deposition Transcript of Nadir Amini (plus exhibits).
- Deposition Transcript of Raj Jaswa (plus exhibits).
- Deposition Transcript of Ali Oztaskin (plus exhibits).
- Deposition Transcript of Bruce Young (plus exhibits).
- Deposition Transcript of Brian Langendorf (plus exhibits).
- Deposition Transcript of George Hayek (plus exhibits).

- Deposition Transcript of Michael Garcia (plus exhibits).
- Deposition Transcript of Christopher Bryant (plus exhibits).
- Deposition Transcript of Brian Reynolds (plus exhibits).
- Deposition Transcript of Michael Collins (plus exhibits).
- Deposition Transcript of Subir Ghosh (plus exhibits).
- Deposition Transcript of Robert Harris (plus exhibits).
- Deposition Transcript of Thomas Shanley (plus exhibits).
- Deposition Transcript of Gary Thome (plus exhibits).
- Deposition Transcript of Ron Epstein (plus exhibits).
- Deposition Transcript of Bernard Marren (plus exhibits).
- Yong Yao, Vendors fight for Pentium core-logic market; Intel is king of the land, most others are losing chip-set share (Cover story), Microprocessor Report (August 21, 1995) (NV 021342-021348).
- Hot Chips VI Symposium Record (SMITH000047-269).
- Viper-M Data Book (OPTINVIDIA 007132-007389).
- Plaintiff Opti Inc.'s Objections and Responses to Defendant NVIDIA's First Set of Interrogatories (EX. 162).
- Plaintiff Opti Inc.'s Objections and Responses to Defendant NVIDIA's Second Set of Interrogatories.
- *Electronic News: Intel readies 'Triton' PCI chipset - Intel's 82430FX - Product Announcement* (Jan. 30, 1995) (NV040393-40394).

12. Throughout this report, I have cited excerpts from depositions and exhibits that have been identified in this case. The citations I provide here are meant to be illustrative, and I may refer to additional portions of the testimony and/or exhibits as necessary and appropriate at trial. In addition, the materials that I plan to use as exhibits in this case will primarily be selected from these materials, although I may later reference exhibits identified by other experts. Additionally, I expect that certain demonstrative evidence, including, but not

limited to, charts, graphs, and drawings, which have not yet been prepared, will be used during my testimony.

13. The charts and timelines that accompany this report are intended to be read together with the body of this report. The inclusion of certain material in the claim charts and/or timelines is not mean to exclude in any way the analysis and opinions set forth in this report. The claim charts and timelines are meant to summarize and augment the opinions that are set forth herein.

V. ANALYSIS OF COURT’S CLAIM CONSTRUCTION RULING

14. I have been given the following instruction regarding means plus function limitations:

7.4 MEANS-PLUS-FUNCTION CLAIM LIMITATIONS

Some patent claim limitations may describe a “means” for performing a function, rather than describing the structure that performs the function. For example, let’s say that a patent describes a table in which the legs are glued to the tabletop. One way to claim the table is to recite the tabletop, four legs and glue between the legs and the tabletop. Another way to claim the table is to recite the tabletop and the legs, but, rather than recite the glue, recite a “means for securing the legs to the tabletop”. This second type of claim limitation is called a “means-plus-function” limitation. It describes a means for performing the function of securing the legs to the tabletop, rather than expressly reciting the glue.


When a claim limitation is in means-plus-function form, it covers the structures described in the patent specification for performing the function stated in the claim, and also any structure that is equivalent to the described structures. In our example, the claim covers a table using glue to secure the legs to the tabletop, as described in the patent, and any equivalent structure that performs the function of securing the legs to the tabletop.

Claims ____ of the ____ patent include means-plus-function limitations. In instructing you about the meaning of a means-plus-function claim limitation, I will tell you, first, the function that each of the means-plus-function claim limitations performs; and second, the structure disclosed in the patent specification that corresponds to each means-plus-function limitation.

[Fed. Cir. Bar Assn. Model J. Instr. (2002)].

15. I have also reviewed the Court’s claim construction for the presnoop patents. Most of the Court’s rulings do not require interpretation for purposes of this invalidity

report. However, the Court's ruling on the "means for initiating a next-line inquiry" requires interpretation. The Court held that the means for initiating a next-line inquiry is "the logic circuitry of the SYSC/IPC schematically illustrated in Figure 9 of the patent."



16. The patent describes Fig. 9 as follows:

FIG. 9 is a schematic diagram of circuitry in the system controller 116 which produces the PSNSTR1 signal used in FIG. 8. As previously mentioned, PSNSTR1 carries a high-going pulse when it is desired to initiate a predictive snoop cycle during a PCI master burst transfer.

Referring to FIG. 9, a three-input NAND gate 902 receives a QPCIFST signal, which is high during the first transfer of a PCI burst or the beginning of a new cache line transfer.

Another input of NAND gate 902 receives a CYCTX signal, which is asserted when both IRDY# and TRDY# are sampled active (a transfer is occurring). NAND gate 902 also receives an LNBREAKB signal, which is low only if the data unit then being transferred is the highest data unit in a cache line. Accordingly, the output of NAND gate 902 will go low during the transfer of the first data unit to be transferred from a line of secondary cache, but not if the transfer is beginning with the highest data unit in the line of secondary memory. This is consistent with the discussion above with respect to FIG. 6 in which predictive snoop is omitted in this situation.

The output of NAND gate 902 is connected to one input of a two-input NAND gate 904, the output of which is connected to the D input of a flip-flop 906. The QN output of flip-flop 906 is connected back to the second input of NAND gate 904. The flip-flop 906 has an inverting clear input which is connected to the output of an AND gate 908, one input of which receives PSNEN, which remains high at all times pertinent herein, and the other input of which receives an EADS1B signal. EADS1B goes low

after EADS#, thereby clearing flip-flop 906. Accordingly, flip-flop 906 latches the output of NAND gate 902 until after EADS# has been asserted.

The Q output of flip-flop 906 is inverted and qualified, in three-input NAND gate 910, by IRDY and MFRAME. IRDY is the inverse of the PCI-bus 118 IRDY# signal, and as previously explained, MFRAME essentially follows the inverse of the PCI-bus FRAME# signal. Thus, NAND gate 910 blocks the output of flip-flop 906 if the PCI device 138 has already indicated that the present transfer is to be the last transfer of the burst. Otherwise, the output of NAND gate 910 (called FTRDTGB ("first TRDY# trigger")) carries a one PCICLK-wide low-going pulse, beginning with the PCICLK rising edge that ends the first PCI transfer of the current line of secondary memory.

The output of NAND gate 910, FTRDTGB, is connected to the D input of a flip-flop 912, which is clocked on LCLKI. Flip-flop 912 thus delays FTRDTGB by one PCICLK to enable other circuitry (not shown) in the system controller 116 to increment the secondary memory line address on HA(31:5) (FIG. 1).

The QN output of flip-flop 912, designated PCIFTRD, is connected to one input of a two-input NAND gate 914, the other input of which receives PSNEN. The output of NAND gate 914 is connected to one input of a two-input NOR gate 916, the other input of which receives the output of another NAND gate 918. One input of NAND gate 918 receives a CSNENDB signal, which is high until EADS# is asserted, and the other input of NAND gate 918 receives the PSNSTR1 signal. The output of NAND gate 916 is connected to the D input of a flip-flop 920 which is clocked by CLK (equivalent to the host bus clock signal HCLK). The QN output is NORed with an inverted version of the Q output of flip-flop 920 to produce the PSNSTR1 signal, which is fed back to NAND gate 918. PSNSTR1 therefor carries a high-going pulse which is synchronized with the host bus clock signal HCLK#, and which remains high until EADS# is asserted.

As previously described, PSNSTR1 is provided to an input of NAND gate 822 in FIG. 8 and, like LT2, initiates an L1 cache inquiry cycle.

[EX. 17, '906 Patent Col. 23:32-24:35].

17. Figure 9 of the '906 patent shows circuitry for asserting the PSNSTR1 signal. The pre-snoop patents explain that "PSNSTR1 will carry a high-going pulse when it is desired to assert EADS# for predictive snoop cycle later in the burst." Thus, the PSNSTR1 signal is initiated in order to initiate the EADS# signal and initiate a next-line inquiry.

18. In order to initiate a next-line inquiry, the secondary memory line address on HA(31:5) must be incremented to point to the next-line. The pre-snoop patents note that the address HA(31:5) must be incremented, but state that the structure to do so is “not shown”:

The output of NAND gate 910, FTRDTGB, is connected to the D input of a flip-flop 912, which is clocked on LCLKI. Flip-flop 912 thus delays FTRDTGB by one PCICLK to enable other circuitry (*not shown*) in the system controller 116 to increment the secondary memory line address on HA(31:5) (FIG. 1).

[EX. 17, ‘906 Patent Col. 24:10-15 (emphasis added)].

19. The PSNSTR1 signal is asserted when IRDY, FRAME, and TRDY are asserted for the first data unit of the current line, the PSNEN signal (presnoop enabled) is asserted as well, and the LNBREAKB signal does not indicate that the transfer is starting on the last data unit of the current line. Thus, the structure in FIG. 9 shows that the PSNSTR1 signal is asserted when the first data unit of the current line is being transferred on the PCI bus (the assertion of TRDY). In particular, the patent indicates that the PSNSTR1 signal depends on the FTRDTGB signal, which is the “first TRDY# trigger.”

Otherwise, the output of NAND gate 910 (called FTRDTGB (“first TRDY# trigger”)) carries a one PCICLK-wide low-going pulse, beginning with the PCICLK rising edge that ends the first PCI transfer of the current line of secondary memory.

[EX. 17, ‘906 Patent Col. 24:5-9].

20. My analysis of the means-plus-function claims for the OPTi presnoop patents, set forth below and in the accompanying charts, will be based on this analysis of FIG. 9 above.

VI. BACKGROUND INFORMATION.

21. I understand that OPTi will present its evidence and witnesses at trial first, before NVIDIA presents its witnesses, and that by the time I testify, the jury should already be familiar with the basic technical terms and concepts in the case. Nevertheless, to the extent necessary, I expect to present summary basic background information regarding the technology,

including how microprocessors work, what a peripheral bus is, how DRAM works, how write-back caches and write-through caches work, what “snooping” is, how cache coherency is maintained, how the PCI Local Bus specification commands work (particularly focusing on MRM), how PCI signals operate (e.g., FRAME, IRDY, TRDY, EADS), how PCI data transfers work (e.g., address phase, data phase), how the PCI Local Bus specification includes limitations on wait states, as well as what prefetching is. Such information will be similar to the information already presented in the tutorials to the Court by both parties, and would be similar to the description of background technology presented in the prior art cited herein. I incorporate NVIDIA’s technical tutorial by reference as a basic summary of the background technology that I may provide.

VII. PERSON OF ORDINARY SKILL IN THE ART.

22. I have been asked to give an opinion regarding the level of skill of a person of ordinary skill in the art. I have been given the following instruction regarding the level of skill in the art:

10.9.3 LEVEL OF ORDINARY SKILL

Obviousness is determined from the perspective of a person of ordinary skill in the art. This person is presumed to know all of the prior art, not just what the inventor may have known. When faced with a problem, this ordinary skilled person is able to apply his or her experience and ability to the problem and also to look to any available prior art to help solve the problem.

Factors to consider in determining the level of ordinary skill in the art include the educational level and experience of people working in the field, the types of problems faced by workers in the art and the solutions found to those problems, and the sophistication of the technology in the field.

[Fed. Cir. Bar Ass’n Model Jury Inst. (2002)].

23. The field of the presnoop patents was relatively narrow and specialized. Relatively few persons were working on chipsets for PCs at the time. Chipset makers included Intel, OPTi, VLSI, UMC, and VIA, and SiS, as well as some PC OEMs, including Compaq and IBM. I understand, from OPTi’s interrogatory response, that OPTi contends that its date of invention for the presnoop patent claims is May 1994. At that time, a person of ordinary skill in

the art in the field of the presnoop patents (PC chipset design) would have had practical experience in chipset design, and in particular, knowledge and understanding of the PCI local bus specification, Rev. 2.0. A person of ordinary skill in the art would have had at least an undergraduate EE degree and some practical experience in chipset design. Due to the fast-moving nature of the chipset design field at the time, and due to the fact that virtually all chipset work was done in industry and in a private setting, academic experience alone would only qualify a person for an entry-level position in the field. Those of ordinary skill in the chipset design art would have had practical experience. The witnesses working in the field who have testified in this case generally had practical working experience.

VIII. INSTRUCTIONS RE: LEGAL STANDARDS

A. Burden of proof instruction.

24. For background on the law, I have been given the following instruction regarding the burden of proof:

1.1 BURDENS OF PROOF

In any legal action, facts must be proved by a required standard of evidence, known as the “burden of proof”. In a patent case such as this, there are two different burdens of proof that are used.

The first burden of proof standard requires that, in order for a party to prevail, you must be persuaded that what the party seeks to prove is more probably true than not true.

The second burden of proof standard is a higher one. It requires that you must be persuaded that it is highly probable that what the party seeks to prove is true.

You may have heard of a burden of proof that is used in criminal cases called “beyond a reasonable doubt”. That requirement is the highest burden of proof. It does not apply to a patent case such as this one, and you should, therefore, put it out of your mind.

I will now give you some background about the nature of this case and the issues you will be deciding. For each issue, I will instruct you as to the burden of proof that will apply. At the end of the trial I will review for you the burden of proof to apply, either the more probable than not standard or the highly probable standard, to each issue in this case.

[Fed. Cir. Bar Ass’n Model Jury Instructions (2002)].

B. Anticipation instructions.

25. I have been asked to give an opinion regarding anticipation by the prior art. I have been given the following instruction regarding the law of anticipation:

10.8 ANTICIPATION/LACK OF NOVELTY

A person cannot obtain a patent on an invention if someone else has already made the same invention. In other words, the invention must be new. If an invention is not new, we say that it was “anticipated” by the prior art. An invention that is “anticipated” by the prior art is not entitled to patent protection. A party challenging the validity of a patent must prove anticipation by the highly probable standard.

In order for a patent claim to be anticipated by the prior art, each and every limitation of the claim must be present within a single item of prior art, whether that prior art is a publication, a prior patent, a prior invention, a prior public use or sale, or some other item of prior art. You may not find that the prior art anticipates a patent claim by combining two or more items of prior art.

A printed publication or patent will not be an anticipation unless it contains a description of the invention covered by the patent claims that is sufficiently detailed to teach a skilled person how to make and use the invention without undue experimentation. That means that a person skilled in the field of the invention reading the printed publication or patent would be able to make and use the invention using only an amount of experimentation that is appropriate for the complexity of the field of the invention and for the level of expertise and knowledge of persons skilled in that field.

In deciding whether or not a single item of prior art anticipates a patent claim, you should consider that which is expressly stated or present in the item of prior art, and also that which is inherently present. Something is inherent in an item of prior art if it is always present in the prior art or always results from the practice of the prior art, and if a skilled person would understand that to be the case.

A prior public use by another may anticipate a patent claim, even if the use was accidental or was not appreciated by the other person. Thus, a prior public use may anticipate an invention even if the user did not intend to use the invention, or even realize he or she had done so.

In this case, defendant contends that claims ____ of the ____ patent are invalid because they are anticipated by _____. If you find that [defendant] has proved that it is highly probable that claims ____ are anticipated, then you must find that the claims are invalid.

C. Prior invention instructions.

26. I have been provided the following summaries of the law of invalidity, which I understand have been taken directly from Model Jury Instructions. The following

description of prior invention was provided to me from the Federal Circuit Bar Association

Model Jury Instructions:

PRIOR ART -- PRIOR INVENTION

An invention made by another person before the inventor on the patent made the invention is prior art to the patent claim, unless that other person abandoned, suppressed or concealed his or her invention.

As a general rule, the first person to reduce an invention to practice is said to be the first inventor. An invention is reduced to practice either when a patent application is filed or when the invention is made and shown to work for its intended purpose. Thus, if another person reduces to practice an invention before the inventor on the patent, then the reduction to practice by the other person will be prior art to the patent claims.

Let's consider an example. Mr. Smith has a patent on a table. He reduced his table to practice on April 1. Ms. Jones invents the same table. She built her table on March 1, one month before Mr. Smith reduced his table to practice. Ms. Jones' invention of the table is prior art to Mr. Smith's patent

claims because Ms. Jones reduced her table to practice one month before Mr. Smith's reduction to practice.

There is, however, an important exception to this general rule. Someone who was first to conceive of an invention but reduced it to practice after someone else will be the first inventor if he or she was the first to conceive of the invention and he or she exercised "reasonable diligence" in reducing the invention to practice from a time just before the other person's conception. Conception of an invention occurs when the inventor has formed the idea of how to make and use every aspect of the patented invention, and all that is required is that it be made, without the need for any further inventive effort. Reasonable diligence means that the inventor worked continuously on reducing the invention to practice. Interruptions necessitated by the everyday problems and obligations of the inventor or those working with him or her do not prevent a finding of diligence.

Let's change our example slightly. Mr. Smith conceived of his table on February 1 and reduced it to practice on April 1. Ms. Jones conceived of the table on January 1, one month before Mr. Smith's conception, and built it on May 1, one month after Mr. Smith's reduction to practice. If Ms. Jones was reasonably diligent in building the table from the time just before Mr. Smith's February 1 conception up to the time that she built the table on May 1, she is the first inventor of the table and her invention is prior art to Mr. Smith's patent claims. The final requirement for a prior invention to be prior art is that the prior inventor did not abandon, suppress or conceal his or her invention. Generally, an invention was not abandoned, suppressed or concealed if the invention was made public, sold or offered for sale, or otherwise used.

27. The following was provided to me from the AIPLA model jury instructions:

Prior Invention

[The defendant] contends that claim[s] _____ of the _____ patent was anticipated because the invention defined in [that claim] [those claims] was invented by another person [X], before [the patentee] invented his invention.

A patent claim is invalid if the invention defined by that claim was invented by another person in the United States, before it was invented by the patentee, and that other person did not abandon, suppress or conceal the invention.

To show anticipation of the [patented] invention, [the defendant] must show by clear and convincing evidence that before [the patentee] invented his invention, [X] actually built a device [practiced a process] [made a composition] that included all of the elements of claim ____ of the ____ patent. In addition, [the defendant] must show that [X]'s device was sufficiently developed that one skilled in the art would have recognized that it would work for its intended purpose. To be an anticipation, all of those acts must have been completed before _____ [the inventor of the patent in suit] invented the invention defined by claim ____ of the ____ patent. [If the prior invention was abandoned, suppressed or concealed, it does not anticipate the ____ patent] However, it is not necessary that [the patentee] had knowledge of that prior invention.

In cases where priority of invention is an issue to be submitted to the jury, further instructions will be required. For example, the jury will need to consider not only the dates when the respective inventions were conceived, but also when the inventions were reduced to practice. An inventor who claims to be the first to conceive of a prior invention but was the last to reduce to practice must also show reasonable diligence from a time just before the other party entered the field until his own reduction to practice in order for the "prior invention" to anticipate the claimed invention in suit.

D. Obviousness instructions.

28. I have been asked to give an opinion regarding obviousness of the presnoop claims in suit. I have been given the following instruction regarding the law of obviousness.

10.9 OBVIOUSNESS

As I mentioned earlier, an inventor is not entitled to a patent if his or her invention would have been obvious to a person of ordinary skill in the field of the invention at the time the invention was made.

Unlike anticipation, obviousness may be shown by considering more than one item of prior art. The question is, would it have been obvious for a skilled person who knew of the prior art to make the claimed invention? If the answer to that question is yes, then the patent claims are invalid. [Defendant] has the burden of proving by the highly probable standard that claims ____ of the ____ patent are invalid for obviousness.

Obviousness is determined from the perspective of a person of ordinary skill in the field of the invention. The issue is not whether the claimed invention would have been obvious to you, to me as a judge, or to a genius in the field of the invention. Rather, the question is whether or not the invention would have been obvious to a person of ordinary skill in the field of the invention.

In deciding obviousness, you must avoid using hindsight; that is, you should not consider what is known today or what was learned from the teachings of the patent. You should not use the patent as a road map for selecting and combining items of prior art. You must put yourself in the place of a person of ordinary skill at the time the invention was made.

You must also keep in mind that the test for obviousness is not whether or not it would have been obvious to try to make the invention, but rather whether or not the invention would have been obvious to a person of ordinary skill in the inventor's field at the time the invention was made.

In determining whether or not these claims would have been obvious, you should make the following determinations:

First, what is the scope and content of the prior art?

Second, what differences, if any, are there between the invention of the claims of the patent and the prior art?

Third, what was the level of ordinary skill in the art at the time the invention was made?

Fourth, are there any objective indications of non-obviousness?

[Fed. Cir. Model Jury Inst. (2002).]

10.9.1 THE SCOPE AND CONTENT OF THE PRIOR ART

Determining the scope and content of the prior art means that you should determine what is disclosed in the prior art relied on by [defendant]. You must decide whether this prior art was reasonably relevant to the particular problem the inventor faced in making the invention covered by the patent claims. Such relevant prior art includes prior art in the field of the invention, and also prior art from other fields that a person of ordinary skill would look to when attempting to solve the problem.

10.9.2 DIFFERENCES BETWEEN THE INVENTION OF THE CLAIMS AND THE PRIOR ART

In determining the differences between the invention covered by the patent claims and the prior art, you should not look at the individual differences in isolation. You must consider the claimed invention as a whole and determine whether or not it would have been obvious in light of all of the prior art.

In deciding whether to combine what is described in various items of prior art, you should keep in mind that there must be some motivation or suggestion for a skilled person to make the combination covered by the patent claims. You should also consider whether or not the prior art “teaches away” from the invention covered by the patent claims. The question to be answered is: Would someone reading the prior art be discouraged from following the path taken by the inventor?

10.9.3 LEVEL OF ORDINARY SKILL

Obviousness is determined from the perspective of a person of ordinary skill in the art. This person is presumed to know all of the prior art, not just what the inventor may have known. When faced with a problem, this ordinary skilled person is able to apply his or her experience and ability to the problem and also to look to any available prior art to help solve the problem.

Factors to consider in determining the level of ordinary skill in the art include the educational level and experience of people working in the field, the types of problems faced by workers in the art and the solutions found to those problems, and the sophistication of the technology in the field.

10.9.4 OBJECTIVE INDICATIONS CONCERNING OBVIOUSNESS

You also must consider what are referred to as objective indications of non-obviousness. Some of these indications of non-obviousness are:

1. Commercial success of products covered by the patent claims or made by a process covered by the patent claims.
2. A long-felt need for the invention.
3. Failed attempts by others to make the invention.
4. Copying of the invention by others in the field.
5. Unexpected results achieved by the invention.
6. Praise of the invention by the infringer or others in the field.
7. The taking of licenses under the patent by others.
8. Expressions of surprise by experts and those skilled in the art at the making of the invention.
9. The patentee proceeded contrary to the accepted wisdom of the prior art.

The presence of any of these objective indications may suggest that the invention was not obvious. These objective indications are only relevant to obviousness if there is a connection, or nexus, between them and the invention covered by the patent claims. For example, commercial success is relevant to obviousness only if the success of the product is related to a feature of the patent claims. If the commercial success is the result of something else, such as innovative marketing, and not to a patented feature, then you should not consider it to be an indication of non-obviousness.

10.9.5 DETERMINATION OF OBVIOUSNESS

[Defendant] contends that the invention claimed in claims ____ of the ____ patent would have been obvious to a person of ordinary skill in the field of the invention at the time the invention was made in light of the ____ item of prior art combined with _____. If you find that [defendant] has proved obviousness by the highly probable standard, then you must find that the claims are invalid for obviousness.

IX. SUMMARY OF OPINIONS.

- A. The asserted claims of the OPTi '906 presnoop patent are invalid in view of the prior invention of Intel employees, including Bruce Young, Ali Oztaskin, and Brian Langendorf (and also published in Shanley).**

B. The asserted claims of the OPTi '906 patent are invalid in view of the prior invention of Compaq employees, including Michael Collins and Gary Thome.

35. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are anticipated under § 102 and/or rendered obvious under § 103 by the prior invention of the snoop ahead method by Compaq employees, including Gary Thome and Michael Collins.

36. The Compaq employees conceived of their invention in late 1992 / early 1993 and reduced it to practice in silicon by February 1994. (OPTi's first conception document is dated May 30, 1994.) The Compaq employees did not abandon, suppress, or conceal their invention, as Compaq filed a patent application in October 1994 (OPTi's filing is July 1995) and released a commercial chipset in November 1994.

C. The asserted claims of the OPTi '906 patent are invalid in view of the prior invention of Motorola employees, including Michael Garcia, Christopher Bryant, and Brian Reynolds.

37. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are anticipated under § 102 and/or rendered obvious under § 103 by the prior invention of the snoop ahead method by Motorola employees, including Michael Garcia, Christopher Bryant, and Brian Reynolds.

38. The Motorola employees conceived of their invention in late 1992 / early 1993 and reduced it to practice in silicon by May 22, 1994. (OPTi's first conception document is dated May 30, 1994.) The Motorola employees did not abandon, suppress, or conceal their invention, as Motorola published several papers regarding the invention and released a commercial chipset in the fall of 1994.

D. The asserted claims of the OPTi presnoop patents are invalid in view of the prior invention of IBM (Microchannel and Amini).

39. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are anticipated under § 102 and/or rendered obvious under § 103 by the prior invention of the snoop ahead method by IBM employees, including Nadir Amini, as set forth in the Amini patent, U.S. Patent No. 5,581,714. [EX. 215].

40. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are rendered obvious under § 103 by IBM's Micro Channel architecture described in the IBM TDB.

X. DETAILED DESCRIPTION OF AND BASIS FOR OPINIONS.

A. The asserted claims of the OPTi patents are invalid in view of the prior invention by Mr. Young, Mr. Oztaskin, Mr. Hayek, and Mr. Langendorf (including the Shanley publication).

192. Incidentally, snoop ahead isn't intended to speed up the uncommon case that the cache contains dirty data. Snoop ahead is supposed to speed up the more common case, where the snoops don't actually hit anything in a cache— when this happens, PCI gets to run at full speed for extended periods, and that's where the performance gain exists.

193. Ghosh “didn't pay much attention” to whether anyone else had addressed the same problem. He testified:

37

18 Q. Is it accurate that you identified the
19 performance bottleneck problem you just mentioned
20 when you were reading the PCI spec?

21 A. There were many specs that I had to go
22 through, including PCI spec.

23 Q. But you identified this problem on your
24 own, no one told you about it?

25 A. That is correct.

38

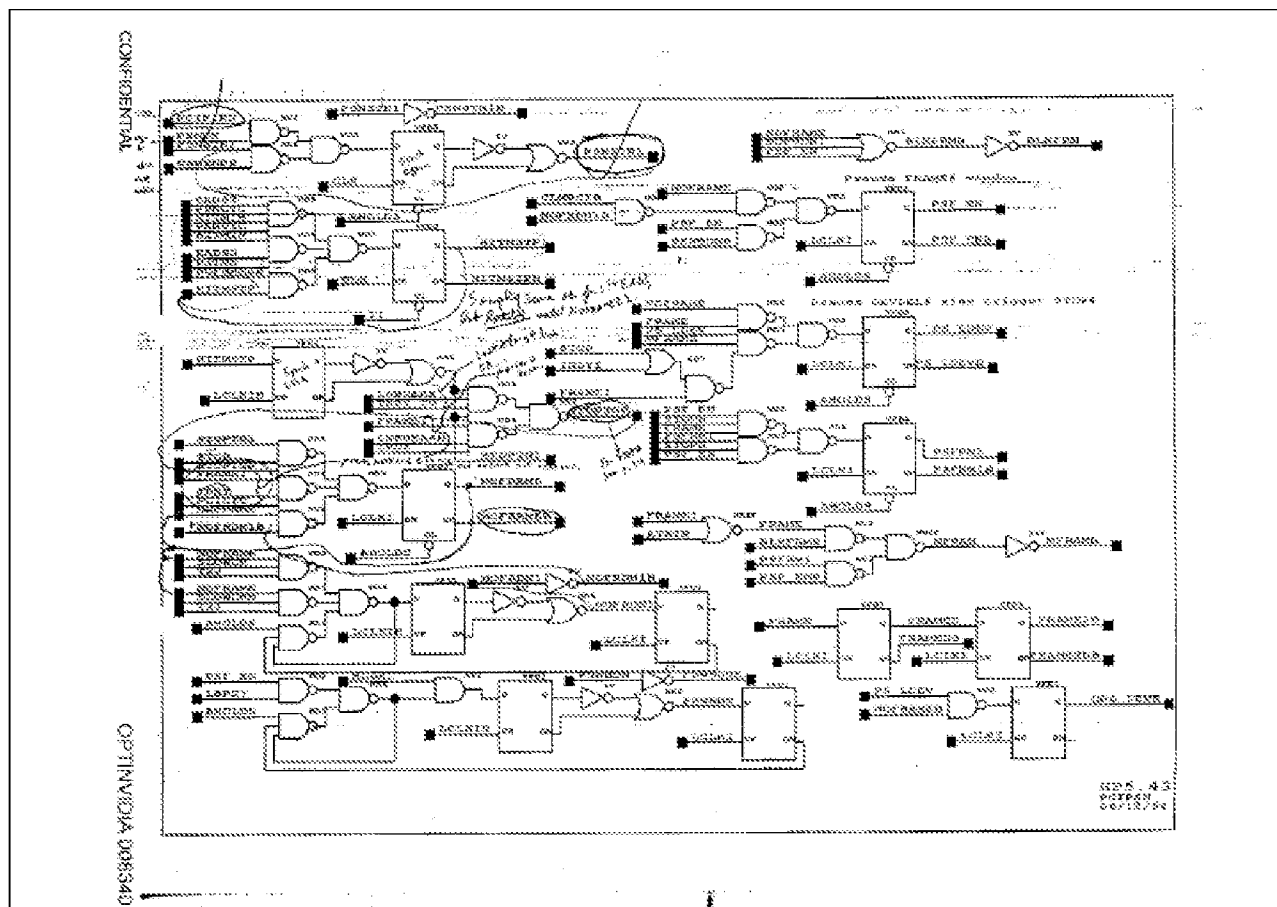
1 Q. At the time, when you identified this
2 performance bottleneck problem in a burst access,
3 did you wonder whether anyone else had identified
4 this same problem?

5 MR. GRACEY: Object to form.

6 You can answer.

7 THE WITNESS: I didn't pay much attention
8 to it.

195. OPTi decided to do the snoop ahead protocol, which Shanley had published earlier. The first dated document showing that OPTi was thinking about presnoop is dated May 30, 1994. [EX. 41, below]. The page below, dated May 30 shows half of FIG. 9 of the '906 patent.



228. If OPTi had invented a chipset that could transfer data units without intervening wait states, it would have disclosed it in the OPTi patent application filed in July 1995, as the “best mode” of practicing presnoop.

B. The asserted claims of the OPTi presnoop patents are invalid in view of the prior invention of Compaq employees, including Michael Collins and Gary Thome.

244. The asserted claims of the OPTi presnoop patent are invalid in view of the prior invention of certain Compaq employees, including Michael Collins and Gary Thome, and described in U.S. Patent No. 5,634,073, filed October 14, 1994 and commercialized in the Triflex/PCI product sold in November 1994.

245. The chipset that Mr. Collins and Mr. Thome invented was called the CMC/DMC chipset. Mr. Collins and Mr. Thome worked at Compaq in Houston, Texas, where they invented the CMC/DMC chipset. Mr. Thome was in charge of the entire project. Mr. Collins was project manager of the North Bridge. Mr. Collins testified:

14

16 Q If you see on the front of Exhibit 182,
17 which is the '073 patent, there are listed five
18 different inventors?

19 A Yes.

20 Q You were the first listed inventor, 1:55:58PM
 21 correct?
 22 A Yes.
 23 Q What was your role in creating the triflex
 24 architecture?
 25 A I was the project manager of this 1:56:09PM
 15
 1 particular north bridge, so of the CMC project I was 1:56:11PM
 2 the project manager.
 3 Q What was Gary Thome's role?
 4 A Gary Thome did, he was in the product
 5 development group so he would be functioning as the 1:56:26PM
 6 total project. So being that this chipset would go
 7 into a product, he would be responsible for the
 8 total project and I was responsible for the
 9 particular north bridge.

246. As cited below and described in detail, Mr. Thome and Mr. Collins both testified that the CMC/DMC chipset both prefetched and presnooped. Mr. Collins noted that in the Compaq system, prefetching and presnooping went together:

55
 16 Q Let me ask this question: Do you
 17 understand what is referred to by the term
 18 "presnoop"?
 19 A Yes.
 20 Q What is your understanding of that term? 3:06:16PM
 21 A Presnooping was, again, it was
 22 speculative. And I'm not sure if it was coming from
 23 the PCI spec or it was in one of the patents, to be
 24 honest with you. Presnooping is when you are
 25 reading ahead, you're speculatively asking for the 3:06:37PM

56
 1 next cache line. That next cache line has to always 3:06:44PM
 2 be snooped because that could be a modified hit in
 3 the caching.
 4 So prefetching and presnooping really go
 5 together, or read ahead. So read multiple, read 3:07:01PM
 6 ahead, prefetching, presnooping, they are all very
 7 similar to the same concept.
 8 Q So did the triflex PCI architecture as
 9 described in the '073 patent implement presnooping?
 10 A I believe so. 3:07:20PM
 11 Q Did it implement it along with
 12 prefetching?
 13 A Yes.

247. OPTi contends that it conceived and reduced the presnoop patents to practice as follows:

Pre-Snoop Patents

(a) U.S. Patent No. 5,710,906 ("the '906 Patent")

The inventions set forth in the asserted claims were conceived sometime prior to May 30, 1994. A working prototype was produced in October of that year and was successfully demonstrated sometime prior to November 12, 1994. Between May 30, 1994 and November 12, 1994 the inventors worked with due diligence to reduce their conceived invention to practice. The first logic diagrams showing the Pre-Snoop circuitry were completed on May 30, 1994. Logic diagrams showing the circuitry disclosed in the patent were completed July 30,

1994. A prototype chip embodying the invention was ordered on August 23, 1994. In the normal course, the prototype would have been received in late October or early November of 1994. The prototype was demonstrated to be operational no later than November 12, 1994. The initial chipsets shipped shortly thereafter. Documents relating to the conception and reduction to practice of the inventions set forth in the asserted claims have been produced with OPTi's P.R. 3-1 Disclosures and Preliminary Infringement Contentions.

[EX. 162, OPTi's Response to NVIDIA's Interrogatory No. 2].

248. Mr. Jaswa, a former OPTi founder who was a VP at the time, testified about the development process at OPTi as follows:

63
9 Do you recall OPTi releasing a chipset
10 in early '95? 12:04
11 A. Yes, yes.
12 Q. And do you recall that that chipset had
13 this presnoop feature included in it?
14 A. Kind of remember, yes.

22 Q. Nine months. So if we wanted to know
23 when development of that chipset started, would we
24 go nine months prior to the date that it was
25 released? 12:04

64

1 A. That's correct, yes.
2 Q. Was that typical for all of OPTi's
3 chipsets?
4 MR. GRACEY: Object to form.
5 THE WITNESS: Yes. As -- we used to 12:05
6 have something called ASICS and we used to have
7 something called, you know, what we call -- okay.
8 We used to have gate array-based technology and then
9 we used to have fully kind of custom, laid-out
10 chips. Custom, laid-out chips, you've got to add 12:05
11 another six months to them. So if it is a gate
12 array implementation, then six to nine months is --
13 six months for development and three months for
14 prototype is the normal development cycle.
15 Q. How would you figure out whether a 12:05
16 chipset was gate array or the other type?
17 A. It depended on kind of gig [sic, "gate"]count. And
18 in those days if the gate count was under -- in the
19 mid '90s if the gate count was under 30,000 or
20 50,000 gates, you could do it as ASIC. Once you got 12:06
21 over 50,000 gates, it had to be a full custom chip.
22 BY MR. TETER:
23 Q. Do you remember which kind the Viper
24 was?
25 A. That was a gate array. 12:06

65

1 Q. So the Viper would have been one that
2 would have had a nine-month development cycle?
3 A. Correct.
4 MR. GRACEY: Object to form.

249. Based on OPTi's interrogatory response, the Compaq inventors, including Mr. Collins, filed their patent application before OPTi reduced its presnoop invention to practice. Filing the Compaq patent application was a "constructive reduction to practice" of the Compaq invention, according to the law as stated in the jury instructions that I cite above.

250. In addition, the testimony shows that the Compaq inventors reduced their invention to practice no later than February 1994, when they received the silicon chips for the CMC/DMC chipset and confirmed that those chipset worked for the PCI specification commands (which would necessarily include MRM), using the PCIFACTs test software. In addition, the silicon was completely prescribed several months earlier, when it was taped out

251. As set forth below, the Compaq inventors also conceived their invention prior to Mr. Ghosh, and diligently worked to reduce that invention to practice. The

Compaq prior invention, and the work reducing it to practice, occurred at Compaq's Houston Texas facilities.

1. Compaq was a member of the PCI special interest group (PCI SIG).

252. As discussed above in the description of Intel's prior invention, Compaq was a member of the PCI special interest group, or PCI SIG. Mr. Collins testified:

101

2 Q Was Compaq part of the PCI special
3 interest group?

4 A Yes, no doubt.

5 Q Was Compaq involved at the very beginning 4:57:33PM
6 of the formation of the PCI special interest group?

7 A I don't know. I think that the initial
8 idea of this type of connection interface bus was
9 probably raised by Intel that Compaq then quickly
10 joined. I do know that we had one of our lead 4:57:59PM
11 architects as the PCI special interest group or SIG
12 group, either the president or chairperson. So we
13 were very active in this spec.

14 Q Would you say that Compaq was on the
15 forefront of PCI design when it originally was 4:58:17PM
16 introduced?

17 A Yeah, no doubt, no doubt, I agree.

253. As OPTi founder Jaswa acknowledged, because Compaq was a member of the PCI SIG, unlike OPTi, Compaq was a generation ahead of OPTi and other chipset makers. Compaq, like Intel, was able to work on PCI chipsets before OPTi.

2. The CMC/DMC architecture included a read prefetch buffer that could hold up to two cache lines in the bridge, allowing the CMC/DMC to burst across cache line boundaries without waiting for snoops at cache line boundaries.

254. The Compaq system was designed to increase performance by having a great deal of "concurrency." In other words, for example, the PCI master could read data from main memory while the processor was working on data in the cache. Mr. Collins testified:

12

24 Q What was the motivation for creating the
25 triflex architecture? 1:53:00PM

13


1 A Performance. 1:53:02PM

4 A So overall system performance, so for
5 example, the reason you have cache controllers on a 1:53:12PM
6 processor is to keep the execution stream going
7 while you're doing some background tasks. So in
8 this case, this has a PCI bus that can read or write
9 the local memory. This has a processor that can
10 operate out of the local cache. And then this also 1:53:43PM
11 has the connection from the processor to the PCI bus
12 directly.

13 So in order to increase performance, the
14 triflex is centered around concurrency, the
15 processor itself can be running locally out of the 1:54:04PM
16 cache and the PCI bus masters can be running locally
17 out of the memory concurrently.

18 The main performance factor is the
19 concurrency, so you can almost be doing two or three
20 items at one time rather than one waiting for the 1:54:28PM
21 other to finish.

255. The Compaq patent illustrated the system in Fig. 1.



256. Figure 4 in the Compaq patent shows the system's North Bridge.




FIG. 4

[EX. 182, '073 Patent].

257. Figure 2 of the Compaq patent showed two squares, or two chips, for the data buffers (212).




FIG. 2

[EX. 182, '073 Patent].

U.S. Patent May 27, 1997 Sheet 2 of 14 5,634,073

258. Mr. Collins, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified that each square in item 212 of Fig. 2 corresponded to Fig. 5:

48

7 Q You referenced earlier with respect to
8 Figure 2 that the data buffers, item 212 had two
9 squares in its illustration?

10 A Yes. 2:53:56PM

11 Q I want to clarify what you meant by the
12 two squares and whether Figure 5 represented a
13 single square in item 212?

14 A On Figure 2 the CMC memory cache
15 controller was a single chip which would be the 2:54:15PM
16 north bridge. The 212 were two separate chips; to
17 be frank it was called the DMCs. Those were two
18 separate DMCs controlled by the CMC. And each one
19 of those chips it was a structure of page 5. So it
20 appears that page 5 is only showing the data path 2:54:43PM
21 through one of the DMCs.

22 MR. BRODY: Figure 5?

23 THE WITNESS: Figure 5. So you can see
24 just by going to the upper left, the PD, that should
25 be 63 through zero. But in this case, it's 63 2:55:02PM


49

1 through 48, which is 16 bits. And then 31 through 2:55:07PM
2 16 which is another 16 bits. And there's a specific
3 reason why that is 16 and 16 are 32, so you would
4 have two chips. What's missing here is 47 through
5 32 and 15 through zero, there being the other chip. 2:55:32PM
6 BY MR. CUNNINGHAM:

7 Q Were there two item 414s, the PCI prefetch
8 buffers?

9 A One in each chip.

259. Fig. 5 shows one of two chips in the Compaq chipset:



3. The CMC/DMC, as described in the Thome/Collins patent, performed a “pipelining” process that both presnooped and prefetched.

260. The Compaq patent describes a chipset for use in systems that supported the PCI bus. The system described in the Compaq patent was designed to work with both the 486 and the Pentium processor. The patent states:

Referring now to FIG. 2, the processor board P is shown. In the processor board P of FIG. 2, the CPU or processor 200 can be any of a plurality of processors, such as the 486DX/33, 486DX2/66, 486DX4/50-100, 486DX4/33-100, 486DX4/33-83, P24T, PENTIUM 50/75, PENTIUM 60/90, and PENTIUM 66/100, and other similar and compatible processors.

[EX. 182, '073 Patent Col. 6:51-57].

261. Mr. Thome testified that the CMC/DMC was designed to work with a Pentium processor, and was the first chipset designed to optimize the performance of second generation Pentium processors.

12

5 Q. The first sentence of the first paragraph in
 6 Exhibit 184 reads, "TriFlex/PCI is the first system
 7 architecture designed to optimize the performance of
 8 PCI-based systems with second-generation Pentium
 9 processors." Do you see that?

10 A. Yes, I do.

11 Q. What were second-generation Pentium processors?
 12 A. The first-generation Pentium processor was
 13 based on -- it was called the Pentium. It ran at 1666 [sic, 66]
 14 megahertz.
 15 And the second one was a -- it was a cost
 16 reduction, it used a different silicone technology, and
 17 was able to run it faster, also, so -- up to 90
 18 megahertz and then beyond. The internal code name for
 19 Intel is called P54C; whereas, the original Pentium was
 20 called P5.
 21 Q. You said "P54C"?
 22 A. Yes.
 23 Q. Was the TriFlex/PCI the first system
 24 architecture designed to optimize the performance of
 25 PCI-based systems with the P54C Pentium processor?
 13
 1 A. To the best of my recollection, yes..

262. Mr. Collins also testified that the CMC/DMC was designed to work
 with a Pentium processor.

26
 19 Q And of course the triflex PCI architecture
 20 was designed to be compatible with the Pentium 2:15:31PM
 21 processors?
 22 A The answer is more than that. Is it all
 23 right? It's designed to work with any processor
 24 that uses a 64-bit data bus with a bursting of 4,
 25 which Intel followed that scheme but others did too. 2:16:03PM

263. Because the CMC/DMC was designed to work with the Pentium, it
 needed to work with a write-back cache. As a result, the CMC/DMC needed to snoop each
 cache line copied from secondary memory. The patent states:

Snooping of read addresses as well as write addresses is necessary
 because the L2 cache controller, and the L1 cache controller in the
 processor 200 in certain cases, are preferably organized as writeback
 cache controllers, and therefore, snoop operations must occur on reads
 to maintain memory coherency.

[EX. 182, '073 Patent Col. 8:47-52].

264. Mr. Thome testified:

17
 5 Q. (BY MR. TETER) Did the CMC/DMC perform snoops
 6 for read accesses by a PCI master?
 7 A. Yes, it did.
 8 Q. Can you explain that?
 9 A. If a PCI master attempts to read some data
 10 location from main memory, then the CMC would -- at the

11 same time that it was reading main memory, would also
12 provide a snoop operation to the second level cache that
13 was controlled by the CMC, as well as the first level
14 cache that was integrated in the P54C processor chip.
15 And so we would provide a snoop operation,
16 inspect it, see if it had -- had the same data [as] main
17 memory, if it was -- then you can leave it alone, or if
18 it had different data, then we would have the cache
19 write back the -- the modified data back to the main
20 memory. And then we would provide that data back to the
21 I/O master on the PCI bus.

265. The patent required each read to be snooped. The patent explained:
Snooping of read addresses as well as write addresses is necessary
because the L2 cache controller, and the L1 cache controller in the
processor 200 in certain cases, are preferably organized as writeback
cache controllers, and therefore, snoop operations must occur on reads
to maintain memory coherency. The PCON block 400 provides the
SNPDONE and SNPHITM or snoop done and snoop hit to modified
signals to the MCON block 404 to allow the MCON block 404 to proceed
with the read or write operations or retry a read operation if appropriate.

[EX. 182, '073 Patent Col 8:47-56].

266. As Mr. Collins testified, every cache line that was transferred to the
read prefetch buffer had to be snooped:

52

5 Q In the read ahead implemented by the 3:00:54PM
6 triflex PCI architecture did every cache line that
7 came into that PCI read prefetch buffer get snooped?
8 A Yes.

267. Because CMC/DMC was designed to work with the PCI bus, it needed
to respond to the "memory read multiple" or MRM command. Mr. Thome testified about his
understanding of the MRM command as follows:

26

12 Q. First of all, what's a Memory Read Multiple
13 command, do you know?
14 A. So the PCI bus specification had two versions
15 of -- of memory reads. The Memory Read Multiple, it was
16 kind of a hint to the memory controller to say, "I'm" --
17 "in addition to making this first request, I'm going to
18 keep on requesting more data in a sequential manner."
19 So the very next address and the address after that.

268. The Compaq patent described a chipset that used prefetch buffers to
perform a memory read multiple, or MRM, command. The patent explains:

The PCI bus provides opportunities to increase overall system performance, particularly that of the memory system. One of the read operations defined for the PCI bus is what is termed as a Memory Read Multiple cycle, which is used to indicate a desire to read a number of cache or memory lines, not just a single line. As noted in the PCI bus cycle definition, this cycle decoding provides an opportunity for the memory controller to start doing read aheads or pipelining so that the data can be obtained prior to actually being required on the PCI bus. With this data then obtained, the memory controller can allow access by the processor, thus further increasing overall system concurrency. However, it is also common for PCI bus masters to abort cycles prior to their completion, and if such an aborted cycle were to occur shortly after a Memory Read Multiple cycle has commenced, then a read ahead operation would have been started and would conventionally complete, only to have the data then immediately discarded. This would reduce overall system performance because of the wasted operations needed to start and complete the full read ahead operation. Therefore, it would be desirable to reduce the wasted time when doing read aheads during Memory Read Multiple cycles when the cycle is aborted early by the bus master on the PCI bus.

[EX. 182, '073 Patent Col. 2:40-63].

269. The Compaq chipset, the CMC, was designed to both prefetch and to presnoop in PCI burst read transfers. The Compaq chipset had a read prefetch buffer, called the "I2M queue", that could hold two cache lines of data. The I2M queue is shown in Figure 5, above. Mr. Collins testified about the I2M prefetch buffer as follows:

46

5 Q Do you understand what I'm referring to 2:49:39PM
6 when I use the term "prefetch buffer"?

7 A Yes.

8 Q What is your understanding of that term?

9 A The prefetch buffer is a buffer in the
10 design that, it's sort of a speculative read, for 2:49:53PM
11 example. So based on the PCI spec, if the PCI spec
12 is issuing a read of memory that asks for read
13 multiple, in some way the second cache line is
14 speculative. You don't know if the master, the PCI
15 master really needs it or not. So it's a way of 2:50:30PM
16 actually doing bursting.

17 Q Did the triflex PCI architecture as
18 described in the '073 patent implement a PCI read
19 prefetch?

20 A Yes. So you actually asked me earlier, so 2:50:44PM
21 Figure 5, so the mechanics are that this is a 64-bit
22 memory data bus that muxes itself down to a 32-bit
23 PCI bus. And it takes two data buffers to do that.

24 This I2M Queue which is number 414, if you
25 look at that, you can see that it's 7 queue 2:51:20PM

47

1 locations deep. So earlier -- 2:51:25PM
2 Q Perhaps 8 if it goes to zero to 7?
3 A Yes, sorry, 8 deep. So you asked earlier
4 a cache line is four times the 64 bits. This buffer
5 being 8 is two times the cache line. So what this 2:51:48PM
6 does is when the initial read of the bus master, the
7 PCI bus master starts, it will load up the first
8 four locations of this buffer. And if it's a read
9 multiple, then when that data is transferred we can
10 transfer the next cache line into this buffer. 2:52:19PM
11 That's why it's actually a depth of 8.
12 The structure is such a way that once you
13 start bursting out the first cache line you can
14 begin filling that buffer with the next cache line,
15 which is a pipelining aspect. 2:52:37PM
16 Q So item 414 in Figure 5 is the structure
17 that would be the PCI read prefetch buffer?
18 A Yes.
19 Q Would you mind circling that on your copy
20 and labeling it as a PCI read prefetch buffer? 2:52:52PM
21 A I should do that?
22 Q Yes, please, so it's clear for the record.
23 A You want a PCI read prefetch.

270. Mr. Collins testified that the prefetch buffers in the system were capable of storing two cache lines:

70

21 Q How many lines can be stored in the I2M
22 Queue?
23 A Two. That's why it's eight locations
24 deep.

71

17 Q So in every implementation of the triflex
18 PCI architecture the PCI read prefetch buffer would
19 be capable of storing two cache lines?
20 A Yes.

271. Mr. Thome also testified that the CMC/DMC buffers could hold two cache lines of data.

25

3 Q. In the portion of the description in the White
4 Paper we were just looking at, it indicated that there
5 is a 64-byte buffer. How many cache lines could be
6 stored in the read buffer in a CMC/DMC?
7 A. Two cache lines.

272. The presence of buffers able to hold more than one cache line supports the testimony of Mr. Thome and Mr. Collins that the Compaq CMC/DMC both presnooped and prefetched. Having space for two cache lines of data in the I2M queue allowed the CMC/DMC to stream across cache line boundaries without waiting for a snoop.

273. Having space for two cache lines, but not using them for presnooping and prefetching, would not make sense. As Mr. Collins testified, the “whole reason” for having prefetch buffers that were two lines long was to allow a PCI read to burst across cache line boundaries.

86

17 Q So the results with respect to the
18 presnooping and prefetching is the that PCI bus
19 master would receive a constant rate of data?

20 A Yes. Well, every clock there will be a 4:23:37PM
21 data transfer, yes.

22 Q Will there with a data transfer across
23 cache line boundaries?

24 A Yes, of course because this is two cache
25 lines, so that's the whole reason the queue is 8.

4. **The CMC/DMC, as reduced to practice in silicon and as described in the Thome/Collins patent, performed a “pipelining” process that both presnooped and prefetched.**

274. Compaq released its CMC/DMC chipset in November 1994.

Q1. What Compaq computers currently utilize TriFlex/PCI?

A1. TriFlex/PCI was first introduced on the ProSignia 500 computer in November 1994. It has since been incorporated into the Deskpro XL Series and the ProLiant 1500 server. It was also incorporated into the new Deskpro and ProLinea Families introduced on March 8, 1995.

[EX. 183 at HP00256 (A Comparison of Compaq Triflex/PCI Architecture to the Intel Triton Chipset at 6)].

275. OPTi has questioned whether the CMC/DMC was able to perform presnooping when it was released in production in November 1994. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified that Compaq released the CMC/DMC chipset in November 1994:

10

7 Q. Do you remember when the CMC/DMC chipset was
8 first used in a commercial product?
9 A. We first started selling it in November of
10 1994.

37

17 Q. And you introduced the CMC/DMC chipset in
18 November?
19 A. November of 1994.

39

15 Q. And underneath it, it reads, "TriFlex/PCI was
16 first introduced on the ProSignia 500 computer in
17 November 1994." Do you see that?
18 A. Yes.
19 Q. Was TriFlex/PCI introduced in the ProSignia 500
20 computer in November 1994?
21 A. So I recall it being somewhere in that time
22 frame. This document seemed to indicate it was November
23 1994.
24 Q. And when we say "TriFlex/PCI," that TriFlex/PCI
25 included the CMC/DMC chipset that we've discussed

40

1 earlier?
2 A. Yes.

136

5 TriFlex/PCI was first released to the
6 public on what date?
7 MR. BRODY: Object to the --
8 A. TriFlex/PCI was implemented in computers at --
9 based on the -- the information in the document, it
10 looks like it was November of 1994. And I recall it
11 being in 1994, as well.

276. Mr. Thome testified that the CMC/DMC chipsets in the Compaq products released between November 1994 and March 1994 were based on the same chipset design, the CMC-2.

113

2 Q. This refers to the products containing the
3 TriFlex chipset, and it gives -- lists three products, I
4 think: the ProSignia product, which launched in
5 November '94, the DeskPro product, DeskPro XL series,
6 and the ProLiant 1500 server, which apparently launched
7 subsequently. And it says it was also incorporated into
8 the DeskPro and ProLinea families, introduced on March
9 8th of 1995, so that, generally, it's a HP 256?

10 A. Uh-huh. I see that.
11 Q. Do you know whether there were any changes in
12 the chipset between those product launches?
13 A. No, there wouldn't have been any changes in the
14 chipset. They were all based on the CMC-2 chipset, that
15 I can recall.

277. Mr. Thome testified that the CMC/DMC chipset was thoroughly tested and debugged and fully functional:

137

7 Q. (BY MR. TETER) Does plaintiff's Exhibit 103
8 indicate to you that Compaq released the TriFlex/PCI to
9 the public in November 1994, without having done the
10 necessary testing?
11 MR. BRODY: Object to the form.
12 A. Compaq would have not -- Compaq would not have
13 released the product until -- until we thought it was
14 thoroughly tested and debugged and fully functional.
15 THE REPORTER: Debugged and?
16 THE WITNESS: Tested and fully functional.
17 MR. DAWSON: Fully functional.
18 Q. (BY MR. TETER) When Compaq released
19 TriFlex/PCI and the CMC/DMC in November 1994, did Compaq
20 have any opinion about whether it was fully debugged and
21 fully functional?
22 MR. BRODY: Object to the form.
23 A. We believe that it was fully debugged and fully
24 functional.

278. Mr. Thome testified that the CMC/DMC chipset had a fully functional MRM command as of November 1994:

141

11 Q. Would you have shipped the CMC/DMC chipset
12 without satisfying yourself that the Memory Read
13 Multiple command would function as you contemplated?
14 A. I don't recall thinking that we would do
15 something like that because without the Memory Read
16 Multiple command operating, we wouldn't be able to run
17 the PCI bus at the full bursting speed that we talked
18 about.
19 Q. And so you wouldn't have shipped product with
20 that feature in it until you had satisfied yourself that
21 it was fully functional and operational, right?
22 A. I think we would have made sure that it was
23 functional before we would ship that product.
24 Q. Okay. Thank you.
25 MR. BRODY: That's all I have.

1 FURTHER EXAMINATION

2 BY MR. TETER:

3 Q. And one last question.

4 There's nothing here today that leads you
5 to believe that it -- that feature wasn't fully
6 functional as of November 1994; is that right?

7 MR. BRODY: Object to the form.

8 A. So to the best of my recollection, that feature
9 was functioning when we shipped the product.

279. OPTi counsel asked whether it was possible that the CMC/DMC was shipped with the MRM command turned off in November 1994. Mr. Thome testified that he did not believe that was true.

15 Q. Is it possible that in November of '94, you
16 shipped product with that feature turned off and that
17 you enabled it later in March when you shipped -- in
18 March of '95, when you shipped further product?

19 MR. TETER: Object to form.

20 A. So to my recollection, we didn't respin the
21 silicone right then. We would have to -- you know, if
22 we could find documentation showing whether we did or
23 not, but I don't recall respinning it, meaning doing
24 another version of the silicone right then.

25 If we -- you know, so -- so if we were

1 using the same silicone all the way through, we wouldn't
2 have had the ability to correct any bugs that existed,
3 you know, from one time to the next.

280. Mr. Collins also testified that the CMC/DMC prefetched and presnooped in response to an MRM command, which is what allowed the system to transfer data on every clock, or achieve the full data rate possible with the PCI bus (with the caveat that another cycle, by the processor for example, did not take priority):

3 Q In normal operation would the triflex PCI
4 architecture presnoop in response to a memory read
5 multiple command? 4:22:44PM

6 A It will try.

7 Q And what effect would this presnooping and
8 prefetching have on the data flow from the
9 perspective of the PCI bus master?

10 A It will be bursting data. It won't be 4:22:58PM

11 interrupted. So, for example, on the PCI bus, every
12 PCI clock there can be a data transfer transaction.
13 If both the master and the slave are enabled with
14 the readies, you can move data with every clock. It
15 would be the goal of the triflex chip to move data 4:23:23PM
16 every clock.

17 Q So the results with respect to the
18 presnooping and prefetching is the that PCI bus
19 master would receive a constant rate of data?

20 A Yes. Well, every clock there will be a 4:23:37PM
21 data transfer, yes.

22 Q Will there with a data transfer across
23 cache line boundaries?

24 A Yes, of course because this is two cache
25 lines, so that's the whole reason the queue is 8. 4:23:53PM

87

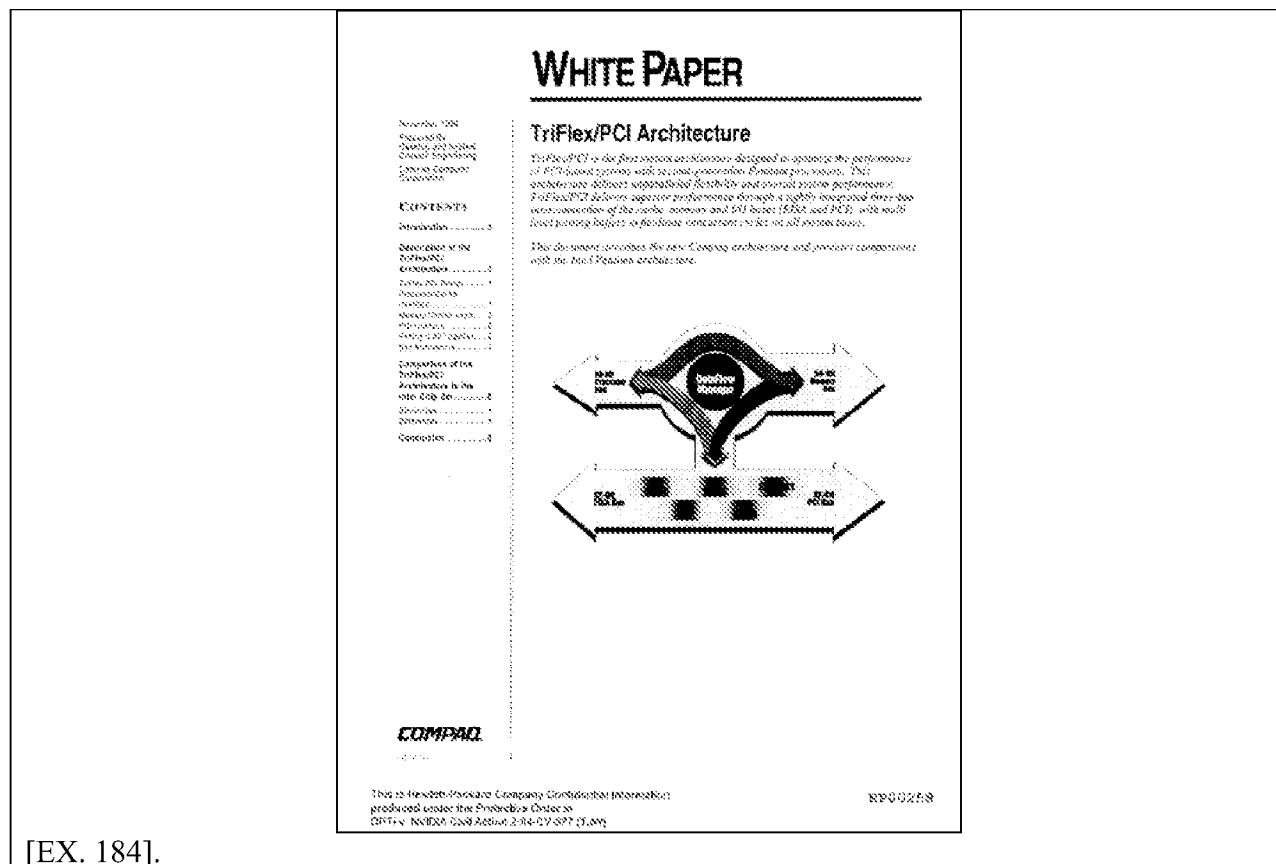
1 Q So as long as the PCI bus master is 4:23:56PM
2 requesting data the presnoop and prefetch will allow
3 that data to be fed to the PCI bus master
4 continuously?

5 A Yes. Well, that is the goal. 4:24:07PM

6 Q In normal operation the PCI bus master
7 will receive the data continuously?

8 A Not quite. In normal operation, if there
9 was no other activity asking the memory to do some
10 other work, if the memory was idle and there was 4:24:26PM
11 constantly read requests coming up or a read
12 multiple -- I don't remember the priority -- but if
13 the PCI read is not conflicting with priority of the
14 memory, then the memory is available, and it will
15 continually do rereads and snoops and keep the 4:24:52PM
16 bursting going of the PCI master, yes.

281. The written evidence, discussed below, supports the testimony of Mr. Collins and Mr. Thome that Compaq had working silicon in November 1994 that had the CMC/DMC capability. Compaq published a whitepaper describing the CMC/DMC during November 1994. Because Intel had not yet released the Triton chipset, Compaq compared the Triflex/PCI, and its CMC/DMC bridge, with the Intel Neptune chipset.



[EX. 184].

282. The operation of the Compaq CMC/DMC chipset is described both in the whitepaper, (EX. 184), and in the Thome/Collins patent, (EX. 182).

283. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified that the CMC/DMC chipset released in November 1994 operated as described in the patent.

80

18 Q. Is it your understanding that the process for
19 handling Memory Read Multiple commands in the CMC/DMC
20 silicon that you received in February 1994 is the same
21 process that you then patented in the 073 patent,
22 Exhibit 182?

23 MR. BRODY: Objection; form. Calls for --

24 Q. (BY MR. TETER) Do you understand the question?

25 MR. BRODY: -- a legal conclusion.

81

1 A. So -- so the way we -- so from CMC-1 to CMC-2,
2 we didn't make any substantial functional changes.
3 Nearly everything we -- we changed were to fix bugs that
4 we found, either in our -- our simulations or in our
5 testing in the lab.

6 And so -- so any changes would have not
7 been from intent in what the design was to do, but
8 because the design, as we developed it in the CMC-1, had
9 some things that it didn't do what we intended to do,
10 such as the errata that we talked about just a few
11 minutes earlier.

81

24 Was it your intent to describe, in Exhibit
25 182, the CMC/DMC chipset that we've been discussing here

82

1 today?

2 MR. BRODY: Object to the form.

3 A. So, yeah. So as -- as I can -- as best I can
4 recall, the intent was for this patent to be a
5 reflection of some aspects of the design of the CMC/DMC.

6 Q. (BY MR. TETER) And focusing on the way the
7 CMC/DMC handled a Memory Read Multiple command, are you
8 aware of or can you recall any differences between the
9 CMC/DMC silicone, that you received in February 1994,
10 and the system described in the patent, Exhibit 182?

11 MR. BRODY: Object to the form. Asked and
12 answered.

13 A. So -- so as -- so as we defined and designed it
14 in February of 1994, the CMC-1 that we received, is
15 essentially the same as what we patented. And any
16 differences would have been, you know, errata for the
17 handful of bugs that we found along the -- along the
18 line.

284. The whitepaper described the CMC/DMC PCI interface as follows:

PCI Interface

TriFlex/PCI is optimized for both processor accesses to the I/O bus and for bus master (EISA and PCI) accesses to main memory on a PCI bus running at 33 MHz.

TriFlex/PCI can send data at 133 MB/s to devices, such as a graphics controller, using the PCI burst protocol.

PCI bus masters may access system memory at the maximum PCI transfer rate, or 32 bits every PCI clock (133 MB/s), improving performance of devices such as the integrated 32-bit Fast SCSI-2 and NetFlex-L ENET controller. Writes from bus masters to system memory are posted in a 64-byte write posting buffer. Reads from bus masters are also stored in the 64-byte buffer, allowing the bus master to read additional data at the 133 MB/s transfer rate. These buffers ensure the bus master maintains the 133 MB/s transfer rate up to 94 percent of all data transfers.

TriFlex/PCI allows EISA and PCI bus masters to access system memory while simultaneously queuing up processor accesses to the I/O bus. This ensures a minimum stall time for the processor and the I/O bus.

[EX. 184 at HP00263].

285. In particular, the Compaq white paper noted that a bus master could read "additional data at the 133 MB/s transfer rate" and stated that the read prefetch buffers

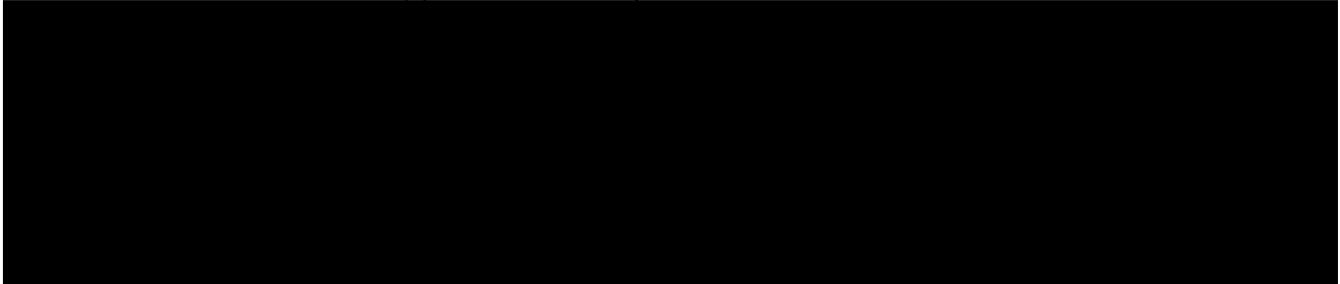
“ensure the bus master maintains the 133 MB/s transfer rate up to 94 percent of all data transfers.”

286. Compaq’s statements about the performance of the CMC/DMC bridge support Mr. Thome’s and Mr. Collins’ testimony that the CMC/DMC chipset used both prefetching and presnooping. CMC/DMC was using prefetching to eliminate DRAM latency between the memory and the bridge, and using presnooping to eliminate the latency that otherwise would occur waiting for a snoop. Thus, the prefetching and presnooping allowed the CMC/DMC to achieve the transfer rates that Compaq recited in the whitepaper.

287. Because the Compaq system both presnooped and prefetched, it allowed the PCI master to burst at the fastest possible rate across the PCI bus. Mr. Collins testified:

88

2 Q So the triflex PCI architecture when
3 implementing the presnoop and the prefetch allows
4 the PCI bus master to burst at the fastest possible
5 rate? 4:25:47PM
6 A Yes, yes. That's the goal.



289. Mr. Collins testified that the Compaq chipset could burst across cache line boundaries as follows:

93

1 Q The result of the triflex PCI architecture 4:33:00PM
2 doing presnooping and prefetching, the PCI bus
3 master can cross cache line boundaries, correct?
4 A Yes.
5 Q Can burst across cache line boundaries? 4:33:11PM
6 A Yes.

290. Recall that Intel, when it was designing Triton, also recognized that to achieve maximum PCI throughput, it should both presnoop and prefetch. If Compaq’s Triflex/PCI bridge prefetched but did not presnoop, it would have had a lower throughput for

PCI burst reads. If Compaq's Triflex/PCI bridge presnooped but did not prefetch, it would have had a lower throughput for PCI burst reads.

291. If either the prefetching or presnooping features of the CMC/DMC chipset was malfunctioning in November 1994, Compaq would not have been able to make the performance claims that it included in the whitepaper. Compaq's performance claims, set forth in the whitepaper, support Mr. Thome's and Mr. Collins' testimony that the CMC/DMC released in 1994 performed both prefetching and presnooping.

292. Comparing the Triflex/PCI bridge with the Neptune, the Compaq white paper noted that the Compaq CMC/DMC bridge achieved a sustained throughput of approximately 125Mbytes/second, compared with Neptune's approximately 75 Mbytes/second throughput.

Differences		
	Compaq Triflex/PCI	Intel Chip Set
Numbers based on:	100/66 MHz Processor	100/66 MHz Processor
DataFlow Manager	Y	N*
Processor to I/O buffer (bits)	256	128
I/O to memory buffer (bits)	512	256
Cache	2-way	direct map
Pipelining supported	Y	Y**
Pipelined memory bus	Y	N
Max. number of cycles simultaneously	16	8
Max. sustained process/memory bus bandwidth (MB/s)	133	92
Sustained I/O bandwidth (MB/s)	123	75

[EX. 184 at HP00265].


293. The Compaq white paper emphasized that Triflex/PCI performed pipelining even when a PCI bus master was reading from secondary memory, unlike the Intel Neptune chipset. The Compaq white paper explained:

TriFlex/PCI is designed with pipelining on the processor bus with EISA and PCI bus master concurrency. TriFlex/PCI always uses pipelining to improve processor performance, even when handling bus master accesses. This is unlike the Intel chip set, which does not allow pipelining when a bus master is running. Further, on the Intel chip set, many pipelined cycles run at the same speed as a non-pipelined cycle. The pipelining support of TriFlex/PCI results in significantly higher sustainable processor/memory and I/O bus bandwidths. Because of the concurrency between the memory pipelining and I/O bus, the memory bus is not left idle when there are cycles that may be run.

[EX. 184 at HP00266].


294. Again, the difference in performance between the CMC/DMC chipset and the Intel Neptune supports Mr. Thome's and Mr. Collins' testimony that the CMC/DMC performed both prefetching and presnooping.

295. The whitepaper contains a description of a process called "pipelining" that illustrates how the presnooping and prefetching worked in the CMC/DMC chipset. The "pipelining" described in the Compaq white paper was a process where one line starts before the data transfer of the previous line is complete. The Compaq white paper illustrated the pipelining process as follows:



[EX. 184 at HP00262].

296. The Compaq white paper contrasted the pipelining performed by the Compaq/ Triflex/PCI chipset with a non-pipelined system, such as Intel's Neptune chipset:



[EX. 184 at HP00262].

297. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified about the CMC/DMC "pipelining"

5 Q. And the second paragraph states, "TriFlex/PCI
6 uses a pipelined memory bus to sustain the maximum
7 throughput of 133 MB/s from the memory. The following
8 diagram illustrates the benefits of pipelining." Do you
9 see that?

10 A. Yes.

11 Q. And then there are two diagrams beneath: one
12 for nonpipeline cycles and one more pipeline cycles?

13 A. Yes.

14 Q. Do you see those?

15 A. Yes.

16 Q. Can you explain what pipelining means, in the
17 context of TriFlex/PCI?

18 A. So pipelining is a -- is a technique to -- to
19 get data to -- to come out faster. So in a typical --
20 the typical way you access memory is, you first present
21 an address to the memory controller and say, "Could you
22 go, please, get the data that's located at this
23 address."

24 And then once -- then the memory controller
25 will go get the data, which takes a certain amount of

19

1 time, and then return the data back.

2 And so if you do that sequentially, then
3 what happens is, you have to wait until the data returns
4 before you can say, "Okay, now here's another address."

5 And since there's a -- as you can see in
6 the picture there, there's an amount of time between
7 when you present the address before you start getting
8 the data returned. You have to wait that time again
9 before you can -- you can go to another cycle.

10 With pipelining, what we do is, we present
11 the address and then the memory controller starts
12 getting the data. And before the memory controller has
13 even returned that first data, we present another
14 address, and say, "Okay, go ahead and start on this one
15 right behind it."

16 And so what that allows us to do is, as
17 soon as the memory controller has returned the first
18 data, it already is -- is working on getting the second
19 data without waiting. And so the net result of that is
20 that we can return data to the requestor at a -- at a
21 much faster pace because we're not waiting for the data
22 to be returned before we start asking for the next
23 address.

298. In the context of a PCI burst read, the pipelining would work as follows. The first line would be fetched by the I2M prefetch queue on the CMC/DMC bridge,

and while the first line was being transferred to the bus master, the second line would be snooped and transferred to the I2M prefetch queue. When the first line had been received by the bus master, the second line would begin to transfer to the bus master without delay for the snoop. At that time, while the second line was being transferred to the bus master, the third line would be snooped, and would then begin to transfer to the buffer. The Compaq pipelining process, would thus use both presnooping and prefetching and would allow the bus master to read data from memory at or very near the maximum rate allowed by the PCI bus.

299. Mr. Thome described the process in the context of a PCI burst read as follows:

25

- 11 Q. Did the CMC/DMC have to stop at cache line
12 boundaries to snoop?
13 A. So we would have to snoop each cache line.
14 Q. Did that slow down the 133 megabyte per second
15 transfer rate of the CMC/DMC?
16 A. Only if the snoop resulted in what's called a
17 snoop hit, where the -- the cache line was modified,
18 data verses what was in main memory.
19 Q. Let's assume that there's a snoop miss, okay?
20 A. Okay.
21 Q. In the case of a snoop miss, would the need to
22 do the snoop slow down the data transfer to a PCI bus
23 master?
24 A. No.
25 Q. Can you explain why not?

26

- 1 A. We would -- once we got a -- a request from a
2 PCI master, we would issue the snoop request to the
3 processor at the same time we would issue the main
4 memory request. And so we could do -- we could almost
5 always snoop the processor faster than the read cycle
6 would take from main memory. And so the snoop would be
7 complete before the memory cycle was complete.
8 Q. If we could sort of walk through one example.
9 So let's say a PCI bus master requests Memory Read
10 Multiple command starting at Line 1, okay?
11 A. Okay.
12 Q. First of all, what's a Memory Read Multiple
13 command, do you know?
14 A. So the PCI bus specification had two versions

15 of -- of memory reads. The Memory Read Multiple, it was
16 kind of a hint to the memory controller to say, "I'm" --
17 "in addition to making this first request, I'm going to
18 keep on requesting more data in a sequential manner."
19 So the very next address and the address after that.
20 Q. So in our example, we'll have a Memory Read
21 Multiple command starting at Line 1, okay?
22 A. Okay.
23 Q. And can you explain the process of requesting
24 data from main memory, transferring data to the PCI
25 master, and snooping, with a Memory Read Multiple

27

1 command that starts on Line 1?

2 A. Okay. So when the -- when the PCI bus issues a
3 Memory Read Multiple request, we would then issue a
4 memory read command to the -- the memory controller, and
5 at the same time, issue a snoop request to the cache
6 controller.

7 Q. Snoop for what line?

8 A. For the same line.

9 So -- so Line 1 -- so we do show a request
10 for Line 1, in the example, to the memory controller and
11 a snoop request for Line 1 to the caches.

12 And then we would let them both run
13 independently, so one would complete before the other.
14 Typically, the snoop would complete first. And if the
15 snoop said that there was no modified data in the
16 caches, then we would just simply wait for the memory
17 read to complete and allow the memory to get buffered to
18 return to the PCI bus.

19 We would also, using the pipelining that we
20 talked about earlier, issue a second memory read
21 command, before the first memory read command completed,
22 of Line 2, to the memory controller, such that we could
23 then continue to -- to read data from the -- from the
24 memory at the faster possible pace.

25 And then when we issued the second read

28

1 command, we would also, at the same time, issue a second
2 snoop command for Line 2, at the same time, as well.

3 And -- and everything would continue going
4 that way, assuming everything the snoop misses, and you
5 would just keep on going.

6 In the case of a snoop hit, we would wait
7 for the memory read to complete from the cache line,
8 because the way memory works, you can't really shut it
9 off in the middle. You have to run it to completion.

10 And then -- then we would write back the
11 data from the processor or from the second cache --
12 second level cache into main memory. Then we would
13 discard the read that we just completed to the PCI bus

14 and issue a new read to reread the data that was
15 discarded, and store that into the -- the buffer, and
16 then present that back out to the PCI bus.
17 Q. In the case of a snoop miss scenario, that you
18 were describing just a moment ago, would the snoop of
19 the next line take place prior to the completion of the
20 transfer to the PCI master of the current line?
21 A. It could.
22 Q. Can you explain the answer -- your answer?
23 A. Yes. Since the PCI master is running at a
24 different speed from the main memory, there is no
25 absolute guarantee or requirement that -- that the

29

1 sequence would happen necessarily one versus the other.
2 Q. In order to achieve the 133 megabyte transfer
3 rate of burst reads to a PCI master, would the snoop of
4 the next line have to occur before the current line has
5 completed transferring all the way to the bus master?
6 A. To the best of my recollection, yes.
7 Q. And is that how it worked in normal operation?
8 A. So to the best of my recollection, yes.

300. Mr. Collins testified that the Compaq system would try to snoop ahead
in normal operation:

85

8 Q. Would the triflex PCI architecture snoop
9 ahead in normal operation?
10 A. It would try. There would be no doubt 4:21:27PM
11 that if the bus master asked for a read multiple and
12 the initial data only occupies half the I2M Queue
13 and there is room for another line, and when that
14 transaction is done, would we go and ask the memory
15 again for the next line and snoop the next line with 4:21:50PM
16 the other half of the I2M Queue empty, and the
17 answer is yes.

301. Mr. Collins' testimony indicates that the trigger for the next line snoop
would be when the I2M prefetch buffer had space for another line. Space in the I2M buffer
would depend on the status of the transfer of the current line. In other words, the next line snoop
would be triggered when the current line began to transfer to the PCI bus master, because that is
when the buffer would have space available for the next cache line. The model instruction for
anticipation in view of the prior art states:

In deciding whether or not a single item of prior art anticipates a patent
claim, you should consider that which is expressly stated or present in the

item of prior art, and also that which is inherently present. Something is inherent in an item of prior art if it is always present in the prior art or always results from the practice of the prior art, and if a skilled person would understand that to be the case.

302. A skilled person would understand that the trigger for the CMC/DMC chipset would be based on the assertion of TRDY and IRDY at the beginning of the current line. Thus, the trigger for a next-line snoop would be the same as the trigger in the '906 patent – as soon as the current line is transferring on the PCI bus, the next line is snooped.

303. The description of the trigger in the '073 patent is consistent with Mr. Collins' testimony:

The ICON block 402 receives this M2INA signal and again checks to see if there is space available in the I2M queue 414 to receive another cache line. When there is room because the ICON block 402 has provided sufficient data to the PCI bus 98, the next M2I read request is provided to the MCON block 404. This process continues until either the Memory Read Multiple completes, a page boundary is crossed or the PCI bus master aborts the cycle.

[EX. 182, '073 Patent Col. 22:29-37].

304. To determine whether space was available during a bursting operation, the CMC/DMC would have to consider the TRDY, FRAME, and IRDY signals on the PCI bus. If the TRDY, FRAME, and IRDY signals were asserted for the first data unit transfer of the line N, the CMC/DMC would know that one cache line in the I2M buffer had become available. In that case, the CMC/DMC trigger would conduct a snoop of line N+1.

305. The whitepaper's description of the CMC/DMC's operation would have been very misleading unless the CMC/DMC chipset had operated as Mr. Collins and Mr. Thome testified in November 1994. There is no reason to believe that Compaq, a publicly-traded company, and its Texas-based engineering team, would have purposely deceived customers by advertising a bursting capability and pipelining feature that did not work for MRM commands. With these features correctly implemented, an easily-verifiable PCI transfer rate could have been measured and compared to the theoretical maximum. Without prefetch/presnoop, maximum measured PCI throughput would have been far lower than the peak. Had their chipset not really implemented the claimed features, the deception would have

been quite obvious. The evidence is clear and convincing that the CMC/DMC chipset, as released in November 1994, implemented presnooping and prefetching, as Mr. Thome and Mr. Collins testified.

306. After Intel released the Triton chipset, Compaq emphasized the Triflex/PCI features. Intel's Triton was capable of bursting across cache line boundaries, but so was Compaq's Triflex/PCI, and Compaq argued that its chipset had a richer feature set. Compaq argued:

PCI to Memory Interface

Intel states that enhancements to their PCI to memory buffers enable bus master transfers from disk to memory to run much faster on Triton than on previous Neptune-based systems, however Triton's >100MB/s peak PCI to memory transfer rate still falls short of TriFlex/PCI. TriFlex/PCI is capable of a peak transfer rate from the PCI bus to system memory of 133MB/s, and capable of sustaining a transfer rate of 125.5MB/s. Further, Triton can only support their high PCI to memory transfer rate as long as the CPU is not accessing the second level cache or system memory. When the CPU accesses the second level cache or system memory in a Triton-based system (which is not an uncommon occurrence), PCI to memory bandwidth falls to zero. Conversely, when a PCI master accesses system memory on Triton, the CPU is prevented from further accesses to either system memory or the second level cache. In contrast, the sophisticated Data Flow Manager of TriFlex/PCI places no such constraints on either bus. With TriFlex/PCI, both PCI bus master and the processor may run simultaneously, largely unimpeded, allowing a higher overall system throughput.

[EX. 183 at HP00251].

5. **The CMC/DMC chipset, as described in the Thome/Collins patent and as reduced to practice in silicon, did not disconnect at cache line boundaries, but limited bursts only to entire pages of memory and/or 256 addresses.**

307. In ordinary operation, because the CMC/DMC allowed MRM transfers to burst across cache line boundaries without disconnecting, the transaction would continue until one of the following events occurred—either the MRM completed, a page boundary is crossed in main memory, or the PCI bus master aborts. [EX. 182, '073 Patent Col. 22:35-37 (“This process continues until either the Memory Read Multiple completes, a page boundary is crossed or the PCI bus master aborts the cycle.”)].

308. The Compaq patent also described a counter that would limit the length of burst read transactions by a PCI master. The counter would allow a PCI master to burst across

many cache lines, but after long bursts of many cache lines, would require the master to relinquish the bus. The '073 patent states:

This M2INA signal is an indication that the processing of the prior address is complete by the MCON block 404 and the processing of the next read address can begin. The actual incrementing of the read address value is performed in the ICON block 402 using an 8 bit counter, thus limiting the total read ahead length to 256 address values. When the counter reaches 255, the read ahead operation is terminated by logic not illustrated for simplicity by causing the Memory Read Multiple to be disconnected. A new address must be received from the PCI bus master to continue the Memory Read Multiple Operation.

[EX. 182, '073 Patent Col. 23:23-34].

309. Mr. Thome, testifying on behalf of HP/Compaq, testified that the counter allowed the MRM burst read to cross up to 255 cache line boundaries in a single burst.

31

6 Q. If you could look, still, at the patent on
7 Column 23 -- Column 23, Line 26?

8 A. Okay.

9 Q. There's a description of, "The actual
10 incrementing of the read address value is performed in
11 the ICON block 402 using 8 bit counter, thus limiting
12 the total read ahead length to 256 address values. When
13 the counter reaches 255, the read ahead operation is
14 terminated by logic not illustrated for simplicity by
15 causing the Memory Read Multiple to be disconnected. A
16 new address must be received from the PCI bus master to
17 continue the Memory Read Multiple Operation."

18 Do you see that?

19 A. Yes, I do.

20 Q. Can you -- well, first of all, was that counter
21 included in the CMC/DMC that was commercialized?

22 A. Yes.

23 Q. Can you explain how that counter worked?

24 A. So when -- when we receive a Read Multiple
25 Command, it -- we were responsible for incrementing the

32

1 address for each additional cache line boundary. The
2 PCI master would just simply give us the first
3 command -- or first addresses and then we had to keep on
4 adding new point -- new addresses.

5 So the counter was used to add -- to
6 increment the address for each -- each additional
7 request. Because it was an 8-bit counter, at some point
8 in time, it would -- you would get to the end, and then
9 at that point, you would have to terminate the PCI
10 transaction and force the PCI master to issue a new --

11 new request with a new address.
 12 Q. So how many cache lines could you cross in one
 13 Memory Read Multiple Operation, using the counter that
 14 you had in the CMC/DMC?
 15 A. So to my recollection, it would be 256 total
 16 cache lines that -- you could cross 255 at the maximum.
 17 Q. Why did you decide to include a counter?
 18 A. If we did not include a counter, then we would
 19 have had to force the PCI master to issue a command -- a
 20 read command for every single cache line that it was
 21 requesting.

33

22 Why did you decide to have a 256 address
 23 value limit for your counter?
 24 A. It seemed like a practical limit.
 25 Q. Why did you decide to have any limit at all for

34

1 the counter?
 2 A. For purposes of not getting too many gates.
 3 As -- as you increase the size of a counter, it -- it
 4 grows dramatically how many gates or transistors are
 5 required to implement the counter.
 6 Q. Was it your understanding that the 255 cache
 7 line limit that you could handle in one MRM transaction
 8 was sufficient for most applications?
 9 A. Our -- our belief was that 256 cache lines
 10 would be sufficient for performance for most
 11 applications.

310. Mr. Collins couldn't recall whether the counter kept track of cache lines or memory addresses, but in any event, testified that the counter allowed a MRM transfer to cross many cache line boundaries:

97

20 Q So when the patent refers to in column 23 4:51:37PM
 21 a total read ahead length of 256 address values, is
 22 it saying that the total read ahead length is 256
 23 cache lines?

24 A That's the part I couldn't find. I don't
 25 know if it's 256 cache lines or it's 256 PCI 4:51:55PM

98

1 addresses. From what I can tell reading the 4:52:02PM
 2 different areas of the patent is this is really 256
 3 addresses based on a PCI bus. So this is 256-bit
 4 addresses of 32 bits.

5 Q Is it accurate to say that the ICON block 4:52:22PM
 6 is keeping track of 256 address values that it is
 7 presenting to the MCON block?

8 MR. BRODY: Object to the form.
9 THE WITNESS: Yes. It's important to
10 note, though, that whatever address, whatever 4:52:39PM
11 address is the initial address by a PCI bus master,
12 that is really the base address. And that is the
13 address that goes to memory to start to fill this
14 queue. Then at every time that the data for that
15 address is returned, you should be incrementing the 4:53:08PM
16 counter.
17 So it's not like ICON has all these
18 addresses, it's basically an incremental address.
19 So for every data transfer you're adding one to the
20 counter. So it's like a running total, 21, 22, 23, 4:53:27PM
21 24, it's an incremental thing up to, this case, 255
22 which is a full counter. So basically you can start
23 anywhere with any address, and then every clock that
24 you do a transfer will increment the counter one
25 notch. And after 256 transfers, that counter will 4:53:48PM
99
1 be at the full rate and that will stop the burst. 4:53:52PM
2 Q And will that always be more than two
3 cache lines?
4 A Oh, yes, much more than two cache lines.
5 Q More than 50 cache lines? 4:54:06PM
6 A Again, is it 256 cache lines or is it 255
7 addresses? So assuming 256 cache lines, it will be
8 256 cache lines, not knowing the fine detail of it
9 if it's 256 addresses, that's a 32-bit address and
10 there are four, so there are eight of those 4:54:33PM
11 addresses to the cache line. So this would be 32
12 cache lines.
13 Q So in any case, regardless of what this
14 means, the PCI bus master will still be able to
15 burst across cache boundaries? 4:54:45PM
16 A No doubt, cache line boundaries.
17 Q Cache line boundaries?
18 A Yes, that's why the I2M is 2X the size of
19 what you need. It's called double buffering. I
20 guess it's important to note that this 256 is kind 4:55:05PM
21 of arbitrary. It could have been a thousand, it
22 could have been 2K. At this time it was, 256 was
23 maybe a recommended off-the-spec or maybe it was a
24 parent data we chose to use but it could be any
25 number.

6. The evidence supports Mr. Thome's and Mr. Collins' testimony that they had working silicon, able to practice the MRM command with presnooping and prefetching, in February 1994.

311. As Compaq introduced the CMC/DMC in commercial products in November 1994, it obviously had reduced the presnooping and prefetching invention to practice in prototype silicon much earlier. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified that Compaq received the prototype CMC/DMC chip, in February 1994. Mr. Thome's notebook supports his contention that Compaq had a working prototype CMC/DMC chip in February 1994.

312. Page HP00039 of Mr. Thome's notebook, EX. 191, dated 2/3/1994, states that "DMC installed . . . no CMC". Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified that Compaq had the DMC chip in silicon as of February 2, 1994:

61

17 Q. If you could look at the page labeled "HP 39."
18 Has a date 2/3/94 in the upper left-hand corner.
19 A. Okay.
20 Q. And there's a figure underneath, and there's
21 some notes to the right of the figure. It says, "CMC
22 CLK. DMC installed. No CMC." And there's something
23 else.
24 Can you explain the status of -- well,
25 first of all, what is that figure?

62

1 A. This appears to be some clock signals measured
2 on the oscilloscope that I spoke of earlier.
3 Q. As of February 3rd, 1994, did you have any
4 silicone for the DMC?
5 A. We did have DMC silicone, based on what this
6 picture and the notes here say.
7 Q. Is that consistent with your recollection, as
8 well?
9 A. It is consistent with my recollection, in terms
10 of roughly where we would have been.
11 Q. As of February 3rd, 1994, did you have any
12 silicone for the CMC yet?
13 A. It would appear that it -- it was not arrived
14 yet. And I will say very, very vaguely, it seems like
15 they had a manufacturing error and they had to go back

16 and re -- redo it, so it wound up taking a little bit
17 longer than originally anticipated.
18 Q. Were the CMC and DMC on two separate chips?
19 A. Yes, they are two separate chips.

313. Page HP00041 is dated 2/17 and states "Got started on FUJI Debug."
Mr. Thome testified that Compaq had the CMC/DMC chipset in silicon by February 17, 1994,
and started on debugging:

62

20 Q. If you could look at the page that has a Bates
21 label HP 41.
22 A. Okay.
23 Q. Is that your handwriting?
24 A. This is my handwriting.
25 Q. And there's a note in the upper left-hand

63

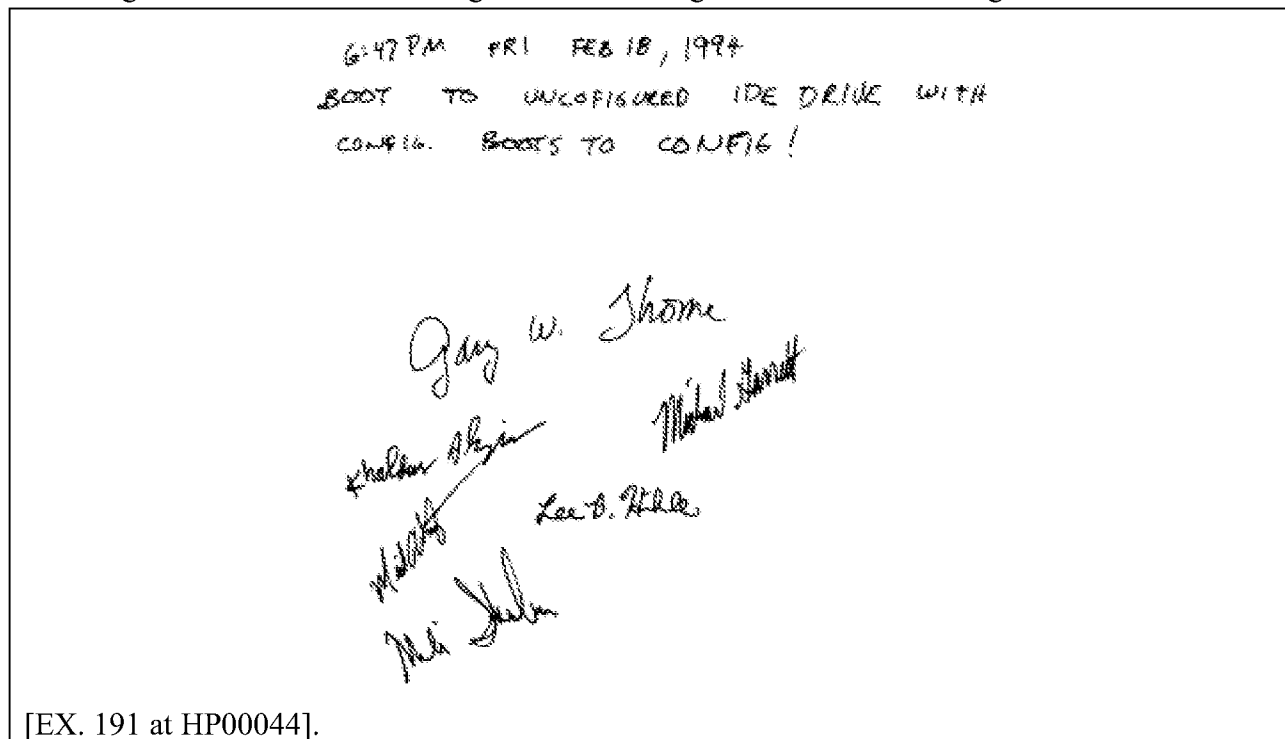
1 corner, "2/17." Does that stand for February 17th,
2 2004?
3 A. It would not be 2004. It would have been --
4 Q. I'm sorry. 19 -- does that note stand for
5 February 17th, 1994?
6 A. I believe this would be February 17th, 1994,
7 because, sequentially, it would have been that year.
8 Q. And there's a time stamp or a time note right
9 next to it, 11:45. Do you see that?
10 A. Yes.
11 Q. And then you write, "Got started on Fugi
12 debug." Do you see that?
13 A. Yes.
14 Q. What does that mean?
15 A. That means that we received all the silicone
16 and manufactured the Fugi board and brought it in the
17 lab and plugged it in and, of course, it didn't work, so
18 we were starting to figure out why it wasn't working
19 correctly.
20 Q. As of February 17th, 1994, did you have
21 silicone for the CMC and DMC?
22 A. I -- I believe so, yes.
23 Yeah, I have notes here saying that we
24 were -- where is it -- the CMC was installed.
25 THE REPORTER: I'm sorry?

64

1 A. I have a note here implying -- implying that
2 the CMC would have been installed.
3 Q. (BY MR. TETER) Can you point out that note?

- 4 A. At the very bottom, it says -- where it says,
5 "Problem: This debug board generates the wrong PC
6 address out of the CMC on the first ROM fetch."
7 So the only way that the CMC could generate
8 the wrong address is if the CMC was installed on the
9 board.
10 Q. So as of February 17, 1994, Compaq, then, had
11 both the CMC and DMC for the CMC/DMC chipset that you
12 described earlier?
13 A. It would appear that way, yes.

314. Page HP00044 is dated 6:47pm, 2/18/1994, and states "Boot to unconfigured IDE drive with config. Boots to config!" with six different signatures.



315. Mr. Thome testified that Compaq had successfully booted an operating system on the CMC/DMC chipset on February 18, 1994:

- 64
- 14 Q. If you could turn to the Page HP 44?
15 A. Okay.
16 Q. And there's a time in the upper left-hand
17 corner, "6:47 p.m., Friday, February 18, 1994." Do you
18 see that?
19 A. I do.
20 Q. And then it says, "Boot to unconfigured IDE
21 drive with config boots to config." Do you see that?

22 A. Yes.
23 Q. And a bunch of signatures underneath. Do you
24 see that?
25 A. Yes.

65

1 Q. What is that note, "boot to unconfig"? What
2 does that mean?
3 A. So it was -- so once we got a first prototype
4 in, it was a bit of a celebration when we could actually
5 get a boot to an operating system. So that meant that
6 it would work sufficiently well enough that it could
7 execute code and get to the point that it could load --
8 like, in this case, it probably would have been
9 Microsoft DOS.
10 And so when we got to that point, we would
11 always -- as soon as it happened, we would write down on
12 our lab notebook that it happened and say exactly what
13 happened, the exact time and then all of us who were in
14 the lab working on it at the time would sign it, just as
15 kind of a -- I guess, a remembrance, more -- more of
16 anything. It was just a happy moment.
17 Q. So as of February 18th, 1994, at 6:47 p.m.,
18 then Compaq had silicone for the CMC and DMC, and it had
19 successfully booted up an operating system?
20 A. That's correct.

316. Mr. Thome testified that over the next several months, Compaq debugged the silicon and worked to make a salable product.

65

21 Q. What did Compaq do after it succeeded in
22 booting up the CMC/DMC with an operating system? What
23 happened over the next couple of months?

24 A. The next several months, we started running a
25 number of different tests on it, and attempted to

66

1 identify any bugs that might have existed in the
2 silicone, corrected those bugs and as well as other
3 bugs, such as some of the software code or firmware
4 code, what we call signal integrity, which is measuring
5 the signals, make sure they're running at the right
6 speed that they're supposed to run at on the board.

7 So there was a number of different tests
8 and procedures that we run over the time to -- to get it
9 fully operational to the point that we could then make
10 it a salable product.

317. Page HP00049 is dated 2/20 and states that "FUJI now boots with L1 and L2 enabled." Mr. Thome testified that Compaq had successfully booted with both caches enabled by February 18, 1994:


66

- 11 Q. If you could turn to the page labeled "HP 49."
12 A. Okay.
13 Q. There's a note at the bottom. It states, "Fuji
14 now boots with L1 and L2 enabled," three exclamation
15 points, remarks after that?
16 A. Yes.
17 Q. And a date and time, "3:15 p.m. 2/20." Do you
18 see that?
19 A. Yes.
20 Q. What does that mean?
21 A. So this means that now -- so when we got it
22 booting, originally, it was booting without the -- the
23 caches enabled. And the caches provided some additional
24 complexity in the system because they -- you know, so
25 having them turned off made it easier to get it running.

67

- 1 What this means is, that we now got it to
2 boot an operating system with the caches turned on.
3 Both -- the L1 stands for the cache that's inside the
4 microprocessor, and the L2 is the cache that's
5 controlled by the CMC chip.
6 So that means that we got it to boot up to
7 an operating system with those caches turned on, and
8 that was a good thing.

318. Page HP00062 states that on February 28, Compaq was running PCIFACTs and looking for bugs.



319. Mr. Thome testified that PCIFACTs was a tool that Compaq used to test a product with every PCI command.

67

20 Q. If you could turn to the page labeled "HP 62."

21 A. Okay.

22 Q. There is a date in the upper left-hand corner
23 of 2/28?

24 A. Yes.

25 Q. See that?

68

1 A. Uh-huh.

2 Q. It says, "PCIFACTs found a CMC bug last Friday.
3 The problem occurs when attempting to run a reserve
4 cycle..." and then it goes on.

5 First of all, is that your handwriting?

6 A. That is my handwriting.

7 Q. What does that note mean, the 2/28 PCIFACTs
8 note that we just read?

9 A. So 2/28 would have been the date, so it was
10 February 28th. That --

11 Q. What year?

12 A. Of 1994, that I made the entry into this.

13 PCIFACTs is a -- is a tool that we used to
14 test out our computers and verify their functionality.
15 And so that -- so my note here was that the PCIFACTs
16 found a combination of cycles that resulted in erroneous
17 operation of the CMC.

18 Q. So in February 1994, you were debugging the CMC
19 and DMC and actual silicone; is that right?

20 A. That's correct, yes.

320. MRM was, of course, a PCI command as of February 1994, and had been so for nearly a year. Compaq, as a member of the PCI SIG, obviously was aware of that, and in August 1993, Compaq was already discussing the high performance that it planned for the MRM command (as compared to the MRL command), which Compaq would achieve by presnooping and prefetching. [EX. 188].

321. Mr. Thome testified that in 1993, Compaq planned to achieve that performance by prefetching and presnooping.

46

11 Q. In August 1993, what did Compaq anticipate the
12 performance of the CMC would be for the Memory Read Line

13 for the 33 megahertz processor?
14 A. I believe it would be 67 megabytes per second.
15 Q. And what did Compaq anticipate at that time
16 would be the performance of the CMC for the Memory Read
17 Multiple command?
18 A. 129 megabytes per second.
19 Q. Why at that time did Compaq anticipate that
20 there would be such a big improvement in the performance
21 for Memory Read Multiple versus the performance for
22 Memory Read Line?
23 A. Because in a Memory Read Multiple command, we
24 could pipeline the -- the memory reads one after
25 another, to continue to sustain the -- the peak

47

1 bandwidth.

2 THE REPORTER: The -- I'm sorry, the what?

3 THE WITNESS: Continue the sustained
4 bandwidth, sustained peak bandwidth.

5 Q. (BY MR. TETER) And is that the process of
6 presnooping and prefetching that you were discussing
7 before?

8 A. It was a process of snooping and a -- and a --
9 and reading ahead the -- the main memory, yes.

10 Q. So as of August 1993, Compaq anticipated
11 performing that process and achieving 129 megabytes per
12 second per Memory Read Multiple command; is that right?

13 A. That is correct.

322. The PCIFACTS tests support Mr. Thome's testimony that the
CMC/DMC chipset was able to run MRM cycles successfully in the February 1994 time frame.

109

8 Q. I didn't see anything in here like that to
9 reflect when you successfully executed the MRM command
10 on the chipset. Is -- is there an entry of that sort in
11 here? And I'm referring to Exhibits 191 and 192.

12 A. Well, so I don't think -- so based on what
13 we've looked at and me skimming through the notes, I
14 don't see anything specifically like that. But if you
15 recall, we talked about something earlier about PCIFACTs
16 running. I did have a note in there about PCIFACTs
17 running?

18 Q. Uh-huh.

19 A. And the way PCIFACTs was designed, it was
20 designed to actually iterate through every single
21 possible PCI command that could be run.

22 And so -- so, you know, to the extent that
23 MRM was, you know, a PCI command, then we would have
24 been implementing it with the PCIFACTs tool and we would
25 have tested it.

1 And so the notebook, I was primarily making
 2 notes for when we found problems. The one exception, as
 3 I mentioned earlier is, yeah, we get excited when we
 4 boot so we, you know, make a note of that --

5 Q. Uh-huh.

6 A. -- kind of a thing.

7 So we -- you know, there's -- you know,
 8 there's PCI, there's PCI read, PCI Read Multiple, PCI
 9 write, you know. And there's a whole bunch of different
 10 commands and we wouldn't make a note every single time
 11 we -- we got one working or not working.

16 Q. If I understand you correctly, then, what
 17 you're telling me is that your recollection is that you
 18 had some software that was running those tests but
 19 there's no written record to confirm that -- that the
 20 tests came back actually showing the MRM command
 21 operational?

22 MR. TETER: Object to form.

23 A. So -- so I think what I'd say is, my
 24 recollection is that we ran the PCIFACTs and, you know,
 25 we would have some notes in here showing, you know, when

1 we were running it and when we found a -- a particular
 2 bug.

3 The way we would do it is, we'd sit down
 4 with the PCIFACTs tool, and first we would run through
 5 every different kind of command. Then we would run
 6 through -- then we'd let it run -- iterate through all
 7 the different variances of that command.

8 So it would like run, you know, a cycle at
 9 offset one byte in the cache and offset two bytes in the
 10 cache, offset three bytes. And so we'd kind of iterate
 11 through all those combinations to make sure that they
 12 all worked properly. And so there were thousands upon
 13 thousands of different possible combinations. And we
 14 would use that software and verify the operation of
 15 the -- of the design that way.

16 And then the meat grinder software that
 17 I -- I spoke of earlier also used the same hardware
 18 cards to generate all the different kinds of PCI bus
 19 cycles, all the different PCI traffic. And that tool
 20 would be one where we would then let them all run
 21 simultaneously and you'd generate a huge amount of
 22 fairly random traffic just to try to find some of the
 23 kind of the -- the corner case bugs, primarily.

323. Mr. Thome's debugging notebooks indicate that he found, and solved,
 bugs in the CMC/DMC chipset.

324. Mr. Thome testified that the CMC/DMC chipset worked as anticipated for the vast majority of cases, and that the bugs appeared only in the very rare combinations of events.

77

3 Q. Now, earlier you testified about how the
4 CMC/DMC would do prefetching and presnooping in a Memory
5 Read Multiple command?

6 A. Yes.

7 Q. You described it as like pipelining, right?

8 A. Uh-huh. Yes.

9 Q. For the basic scenario where a PCI master is
10 just reading from secondary memory and there aren't any
11 other requests going on by the processor or anyone else,
12 did the CMC/DMC silicone that you received in February
13 1994 work for that basic process?

14 MR. BRODY: Object to the form.

15 Q. (BY MR. TETER) Do you understand the question?

16 A. I do understand the question.

17 So -- so I guess what I would say is that
18 the -- the CMC that we received worked correctly for
19 nearly every combination. Bear in mind, there were
20 probably millions of different possible combi --
21 combinations, and there were -- the ones that didn't
22 work were probably in the -- in the tens or maybe
23 hundreds, at most, of combinations.

24 So nearly everything did work correctly
25 as -- as we expected or anticipated it to.

78

1 Q. This bug that you identified, is this what you
2 would sometimes describe as a "corner case"?

3 A. Yeah. So we would say it was an unusual, you
4 know, sequence, and it took an awful lot of effort to
5 even -- to get it to occur. And so it was one that
6 was -- you know, you could -- you could describe as a
7 corner case, although it was certainly important to have
8 it fixed before we took it to production because,
9 otherwise, you would -- you know, even -- even though it
10 only happened infrequently, infrequently having the
11 wrong data is not a good thing for customer
12 applications.

13 Q. So for most case, the CMC/DMC was able to
14 perform the prefetching and presnooping, that you
15 described before, at least as it -- as you had the
16 silicone in early 1994; is that right?

17 A. Yeah.

18 MR. BRODY: Object to the form.

19 A. Okay. So the fact that it took us three months

20 before we saw the bug, it was clearly operating
21 correctly for most cases..

325. The bugs that Compaq identified and solved in the CMC/DMC chipset were for what I have described above as "corner cases" in a design. The entries in the Compaq notebook do not indicate any problem with the PCI MRM command, which tends to indicate that the command was working as anticipated. If the PCI MRM command had *not* been functioning, one would have expected to see many entries about it. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified:

133

17 Q. You mention that there was a -- a program that
18 you had called PCIFACTs. Do you remember mentioning
19 that?
20 A. Yes.
21 Q. If you look at Page HP 62 of Exhibit 192?
22 A. Okay. I'm here.
23 Q. And you have the note, "2/28, PCIFACTs found a
24 CMC bug last Friday." Do you see that?
25 A. Yes.

134

1 Q. Does that indicate that you were running
2 PCIFACTs on the silicone in February 1994?
3 A. Yes, it does.
4 Q. Now, if you had had or found a bug -- strike
5 that.
6 Did -- did PCIFACTs run tests on all PCI
7 commands?
8 A. Yes, it did.
9 Q. In 1994, did PCIFACTs then run a test on MRM?
10 MR. BRODY: Object to the form.
11 A. I believe that we'd run tests for MRM. And I'm
12 just looking at the timing diagram right here -- and I
13 have to refresh my memory -- but it may actually even be
14 a MRM command that's on the timing diagram right here.
15 Q. (BY MR. TETER) If you had found a bug, in
16 February 1994, on the silicone for the MRM command,
17 would you have recorded it in your engineering notebook?
18 A. Likely. Probably would have, but I attempted
19 to record the bugs as we found them, in the notebook.
20 Q. And the notebook doesn't indicate any
21 particular bug for MRM, does it, not that you saw in
22 your review last night?
23 A. I don't recall seeing anything particular call
24 out a MRM.
25 Q. Is it consistent with your recollection that

- 1 the MRM command in the silicone that you received in
 2 1994 February operated as expected?
 3 MR. BRODY: Object to the form.
 4 Q. (BY MR. TETER) Do you understand the question?
 5 A. Yes. So my recollection is that it generally
 6 operated correctly with the MRM command.
 7 Q. And when we say "operated correctly," are you
 8 referring to doing the prefetching and presnooping that
 9 you described earlier on direct?
 10 MR. BRODY: Object to the form.
 11 A. So my recollection is that it would have done
 12 the presnooping and prefetching.
 13 Q. (BY MR. TETER) As of February 1994?
 14 A. As of February of 1994, yes.

326. Thus, the written evidence, including the Compaq lab notebooks and whitepaper, support the testimony of Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, that Compaq had reduced the presnooping and prefetching invention to practice in February 1994.

327. In 1994, after debugging the silicon for the prototype, Compaq had the second version of the silicon, CMC-2, manufactured.

- 21 Q. If you could look at the bottom line of Page
 22 HP 62?
 23 A. Yes.
 24 Q. It reads, "No CMC-1 workaround; will be fixed
 25 in CMC-2." Do you see that?

- 1 A. Yes.
 2 Q. What was the difference between CMC-1 and
 3 CMC-2?
 4 A. Well, CMC-1 was the -- the first version of the
 5 chip that we used to bring in the lab and start
 6 debugging. And as we found bugs in it or errata, we
 7 would make those corrections and then we'd manufacture
 8 the second version of -- of the CMC chip, which was
 9 called CMC-2. And that second version is the one that
 10 we ultimately took to production, if I recall.
 11 Q. Now, when you talk about these two different
 12 versions, CMC-1 and CMC-2, does that correspond to what
 13 you were describing before as -1 or A0 samples in -2 or
 14 A1 samples, in Exhibit 188?

15 A. Yes. So we always use the terminology -1 and
16 -2 for our versions.

18 Q. When did Compaq receive CMC-2?

19 A. It would have been sometime in 1994, and I
20 don't recall the exact date.

328. Supported by his lab notebooks, Mr. Thome testified that Compaq received the CMC-2 silicon sometime in August 1994.

107

21 Q. Okay. What happened between May 24th and
22 August 30th, do you know, on the CMC/DMC?

23 A. Well, so I -- I don't recall specifically, but
24 I will say that probably what happened here is, this is
25 probably we -- this was probably the last bug that we

108

1 found in the CMC-1. And sometime a little bit after
2 that, we released CMC-2. And as I mentioned, it
3 seemed -- you know, it takes several weeks, you know, a
4 couple months to -- once you release the silicone before
5 you actually get it back.

6 So if I were to look at this, I'd say we
7 probably released CMC-2 probably sometime in -- in June,
8 and then started receiving the silicone back in August,
9 started building up more boards.

10 And then the -- the notes on August 30th
11 is -- is reflective of what we called our wide area
12 test, where we go from kind of our engineering lab, the
13 engineering development lab and start moving it into a
14 number of other test labs where they test a lot of -- a
15 much broader set of -- of third-party adapter cards,
16 third-party software, and so forth like that.

17 And this -- that's the kind of work I was
18 doing at the end of August.

19 Q. Okay.

20 A. That's -- that's a little bit of speculation on
21 my part. But it looks to me, based on the information
22 and the time -- timeline that would probably have been
23 roughly accurate.

329. Mr. Thome testified that it took Compaq approximately 2 months after determining that a chipset was ready for production to include it in a commercial product.

44

5 From the time that a chipset is ready to be
6 produced, to the time it actually ends up in products,
7 what's required -- or what was required at Compaq?

8 A. We would get the -- you know, the second
9 version of the prototype in or whatever the production
10 version of the -- the silicone in, and we would test it
11 for -- for a -- usually some amount of time and it
12 would -- depending on how substantial the changes were,
13 could range from a couple weeks to a few months.
14 And once we felt confident that it was
15 going to work, we would start the -- we would place an
16 order to our production fabrication facility that would
17 make the production versions of the chips, and that
18 would usually take several weeks to, say, maybe about --
19 maybe about six weeks before they get the first version
20 for those completed. And then it would take us another,
21 roughly, two weeks to put them on motherboards, put them
22 into computers, and then start shipping the product to
23 customers.

330. Mr. Thome's timeline is completely consistent with the launch of the CMC/DMC in November 1994 indicated by the Compaq/HP documents.

7. Mr. Thome and Mr. Collins conceived of the presnooping and prefetching invention in the 1992-93 time frame, before OPTi.

331. Mr. Collins and Mr. Thome testified that they conceived of the presnooping and prefetching invention early in the design of the CMC/DMC chipset.

332. Mr. Thome, on behalf of Hewlett-Packard/Compaq as a designated corporate representative, testified:

8

24 Q. When did you start working on the CMC/DMC?
25 A. It was in 1992.

53

9 Q. (BY MR. TETER) By what date had Compaq decided
10 to have an I2M queue buffer that would store two cache
11 lines of data?

12 A. I don't recall the exact date, but it would
13 have been either sometime in the second half of 1992, or
14 the -- probably the -- or possibly the first half of
15 1993.

16 Q. By what date had Compaq decided to do the
17 prefetching and presnooping, that you described earlier,
18 in the context of a Memory Read Multiple command?

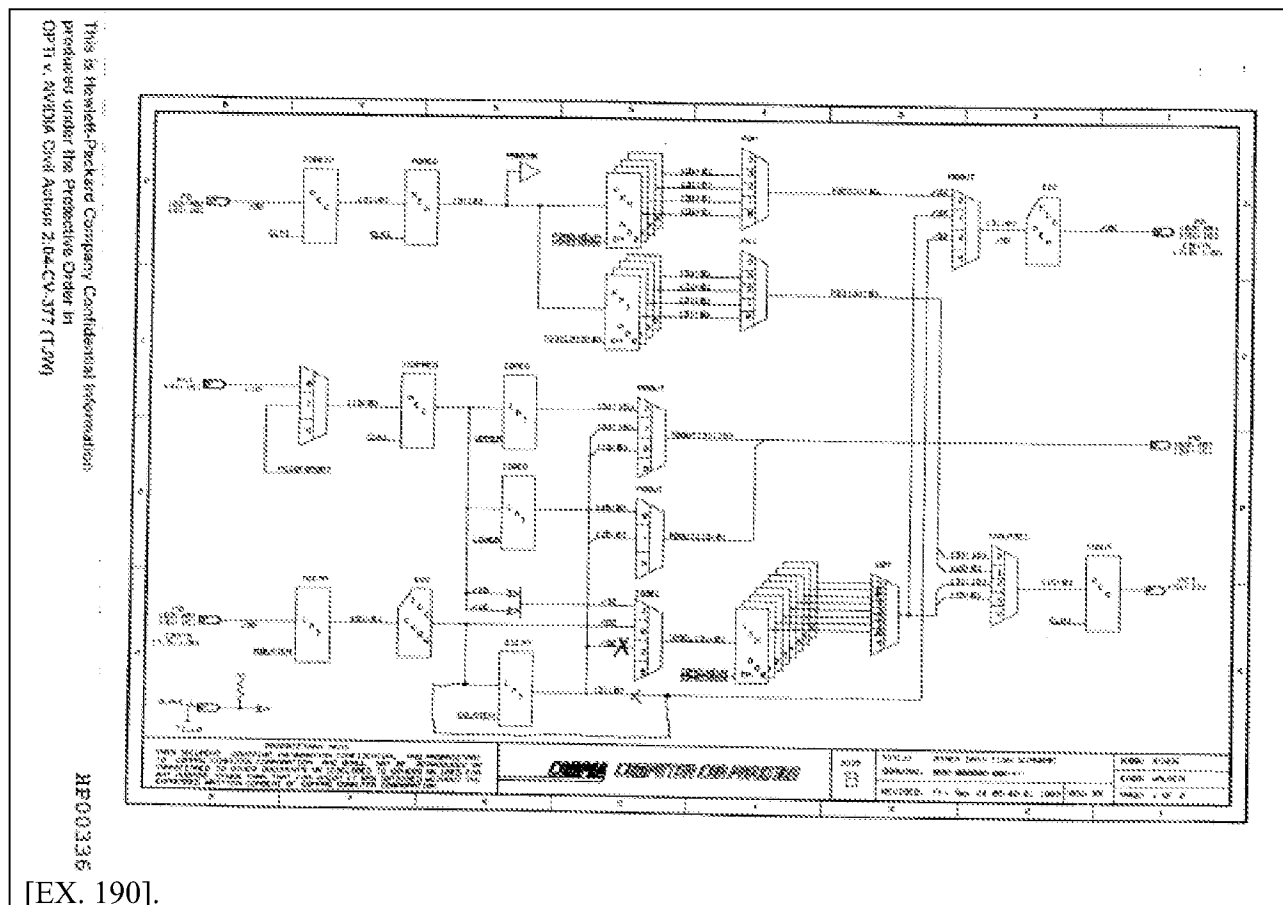
19 MR. BRODY: Object to the form.

20 Q. (BY MR. TETER) Do you understand the question?
21 A. I do understand what the question is.

21 data transfer, yes.
 22 Q Will there with a data transfer across
 23 cache line boundaries?
 24 A Yes, of course because this is two cache
 25 lines, so that's the whole reason the queue is 8. 4:23:53PM

337. Mr. Thome testified that Compaq already had conceived of the presnooping and prefetching method by August 1993, the date of EX. 188.

338. In addition, a Compaq drawing of the Biner chipset (which used the same architecture as CMC/DMC) dated September 24, 1993, shows the Compaq I2M queue, and illustrates that it can hold two cache lines of data:



339. As noted above, having buffers that are able to hold two cache lines of data strongly supports Mr. Thome's and Mr. Collins' testimony that they conceived of the presnoop and prefetching idea from the very early stages of the process. In fact, Figure 5 of the patent is virtually identical to the September 1993 Compaq drawing.

340. Furthermore, in order to have silicon in February, the Compaq engineers would have completed the design, and taped it out, earlier.

8. The documents that OPTi cites do not contradict Compaq's conception.

341. During the depositions of the Compaq inventors, OPTi's counsel argued that typewritten documents dated in late 1992 / early 1993 contradicted the testimony of Mr. Collins and Mr. Thome. According to the argument of OPTi counsel, the typewritten documents showed that Compaq planned to disconnect at cache line boundaries and that Compaq did not plan to have a streaming feature.

342. Neither Mr. Collins nor Mr. Thome knew who wrote the typewritten documents or what exactly they represented. Mr. Collins speculated that Michael Moriarty might have written the document, but nevertheless remained confident that the Triflex/PCI architecture permitted PCI bus masters to burst across cache lines by 1993:

188

17 Q Are you confident that -- do you know who
18 authored Plaintiff's Exhibit 100?

19 A This could be Michael Moriarty.

20 Q Why do you believe that Michael Moriarty 9:08:42PM
21 might have authored Plaintiff's Exhibit 100?

22 A Because these are very specific to the
23 memory, excuse me, to the PCI reads and writes.

24 Q This is dated in 1992, correct?

25 A Yes. 9:09:05PM

189

1 Q Are you confident that as of 1993 the 9:09:06PM

2 triflex PCI architecture permitted PCI bus masters
3 to burst across cache line boundaries?

4 A Yes.

343. Responding to questions from OPTi's counsel, Mr. Thome testified that he recalled always wanting to do read ahead (prefetching and presnooping) from the start, and emphasized that it was included in the February 1994 silicon:

103

4 Q. And the reason I wanted to point you to that
5 passage was to ask you whether that indicates that the
6 plan, in December of '92 at least, was to do bursts only
7 of one cache line at a time?
8 A. So in my recollection, I don't recall ever
9 making a decision that we would do that. The -- the
10 passage here would seem to indicate that that was what
11 our thinking was at the time or --
12 Q. Okay.
13 A. -- whoever wrote the document, that was their
14 thinking at the time. I -- I was not the -- the author
15 of the document.
16 Q. Okay.
17 A.[sic, Q] But your recollection is that, eventually, you
18 changed that, and you decided to burst more than one
19 cache line at a time?
20 MR. TETER: Object to form.
21 A. So -- so my recollection is, as long as I can
22 remember discussing it, we had decided to -- to burst
23 it. But...
24 Q. (BY MR. BRODY) Burst more than one cache line?
25 A. Burst more than one cache line on read -- on

104

1 Memory Read Multiple commands.
2 But as I said, you know, I was not the only
3 person on the -- on the design, so it's certainly
4 possible that either we were all thinking this at -- at
5 the time or some people were thinking at the time and
6 not all of us.
7 Q. Do you know when the group came to consensus
8 that you were going to burst more than one cache line?
9 Can you give me a date?
10 MR. TETER: Object to form.
11 A. Well, so it would have been before we re --
12 received the CMC-1, because it was part of the feature
13 set from --
14 Q. (BY MR. BRODY) Okay.
15 A. -- from CMC-1.
16 So it would have been sometime relatively
17 early in the design cycle of the CMC-1.
18 Q. Okay. You got the CMC-1, and I think we were
19 looking at February or March of '94?
20 A. Yes.

344. Mr. Collins also disagreed with the interpretation of the documents by OPTi counsel, despite a lengthy examination. Mr. Collins testified that he believed that the documents were describing bursting in another context and that the document did not affect his

belief that, as of 1993, the Triflex/PCI architecture permitted PCI bus masters to burst across cache line boundaries.

170

19 Q Isn't that possible?

20 A Anything is possible, but I cannot imagine 8:28:22PM
21 that we would be patenting a concept that's not in
22 the chip.

23 Q The fact that you ultimately got sometime
24 in 1994 to multiple lines of buffers, multiple line
25 buffers, I should say, and that that's, may well 8:28:39PM

171

1 have gone into the chip and may have gone into your 8:28:43PM
2 patent that it's not inconsistent with the fact that
3 earlier on you may have had a different architecture
4 in mind?

5 MR. CUNNINGHAM: Object to form. 8:28:54PM

6 THE WITNESS: Again, I'm disagreeing with
7 you because I think we are reading the way this is
8 as bursting as not being the meaning that we were
9 doing. I think you're trying to -- it's the PCI
10 burst and I think it is really the fact that the 8:29:06PM
11 lead-off address of a PCI read would be always
12 reading one line or not.

13 So, for example, this could be a memory
14 burst that a PCI is one 32-bit address, but the
15 memory system out back is eight times that. So we 8:29:27PM
16 would be doing the bursting but not of more than one
17 line which would be eight times the data of what the
18 PCI's asking for. I think that's where I disagree.
19 It's not the PCI burst.

186

17 Q If I could direct you to Exhibit 168 which
18 is the schematic that is Bates labeled HP336. Are
19 you confident that the triflex PCI architecture
20 permitted PCI bus masters to burst across cache line 9:05:03PM
21 boundaries?

22 A Yes.

23 Q And are you confident that Exhibit 186
24 illustrates the implementation of the architecture
25 that permitted PCI bus masters to burst across cache 9:05:21PM

187

1 line boundaries? 9:05:25PM

2 A From a data path perspective, the data
3 path is here for two cache lines to be double
4 buffering.

5 Q You're confident as of September 24th, 9:05:39PM
6 1993, Compaq had implemented a system that allowed
7 PCI bus masters to burst across cache line
8 boundaries?
9 MR. BRODY: Object to form.
10 THE WITNESS: Again, not recalling the 9:05:57PM
11 exact detail, I believe so. I mean this is on the
12 order of maybe almost a year after this project.
13 And like I said earlier, 183 is a follow-on to the
14 DMC. This is the Biner chips. So this is the
15 second generation data controllers for ECC. 9:06:18PM
16 BY MR. CUNNINGHAM:
17 Q You have no reason to believe that the
18 date on Exhibit 186 is inaccurate, do you?
19 A No, I don't.

345. Mr. Collins only caveat was that the Triflex/PCI architecture, in 1993, might have had a restriction on bursting across cache lines for a write-through (not write back) cache—which is a different context that does not involve snooping and is not relevant:

189

7 Q And as of 1993 the triflex PCI
8 architecture would presnoop and prefetch in response
9 to a PCI memory read multiple command?
10 MR. BRODY: Object to form. 9:09:28PM
11 THE WITNESS: To answer, with limitation.
12 I think what was pointed out earlier there seems to
13 be some restriction on write-through address tags.
14 Now I've lost my hundred percent confidence.

346. Even if the Triflex/PCI architecture did have a restriction on bursting in the write-through context in 1993, it would still have burst across cache lines when operating with a Pentium microprocessor because the Pentium microprocessor utilized a write-back cache. Thus, Mr. Collins remained confident even at the very end of his deposition that the Triflex/PCI architecture permitted PCI bus masters to burst across cache line boundaries:

189

1 Q Are you confident that as of 1993 the 9:09:06PM
2 triflex PCI architecture permitted PCI bus masters
3 to burst across cache line boundaries?

347. Mr. Thome testified that the product as shipped did the read-ahead (presnooping and prefetching) for all PCI MRMs.

5 Q. Do you recall that in December of 1992, you
6 were going to prefetch for some PCI to memory reads but
7 not for others?

8 A. So -- so, yes, we wrote this down, presuming
9 this is what we were doing. I don't recall what -- you
10 know, what the distinctions were, in terms of when we
11 would do the prefetching and when we would not.

12 Q. Okay. Do you know that -- if that's the way
13 you ultimately did it?

14 A. The way we ultimately did it -- so the
15 configuration registers were used to be able to --
16 primarily to do things like either set a -- set a
17 setting for something or to be able to turn features on
18 or off.

19 Q. Uh-huh.

20 A. So -- and as I can vaguely recall, I believe
21 that we probably had a configuration registered to be
22 able to turn off the feature of being able to -- to read
23 multiple cycles or, you know, read the -- you know,
24 prefetch and read multiple cycles.

25 And so I'm -- I'm kind of saying this from

96

1 really a vague recollection because I don't recall the
2 specifics of -- of how we implemented exactly. But we
3 probably had something to turn that off because that
4 would have been a feature we would have been concerned
5 about possibly having bugs in. And so we said, well, we
6 can turn it off so that we can get around it, if we need
7 to, to continue debugging other functions of the -- of
8 the chipset.

9 Though, when we shipped it, we shipped it
10 with PCI Read Multiples always reading ahead.

11 Q. Okay.

12 A. And the PCI standard read cycle, we may have
13 had a configuration register option to decide whether or
14 not to prefetch or not prefetch.

15 As I -- as I recall, I think our concern
16 was it was unclear for -- so it was unclear because the
17 PCI specification was new at the time. And so it was
18 unclear if -- if PCI bus master devices would issue PCI
19 Read Multiple commands, when what they really meant
20 to -- or issue PCI read commands when what they really
21 meant to do is issue PCI Read Multiple -- Multiple
22 commands.

23 This is coming back to me slowly here. So
24 I think what we did is, we had a register to say if we
25 found a PCI device that was running PCI read commands

97

1 when they really -- what they really wanted was to issue

2 PCI Read Multiple commands, we would treat that PCI read
3 command as if it was a PCI Read Multiple command and run
4 the prefetch cycles for that device on -- on its behalf.
5 In effect, treat it like it was a Read Multiple command,
6 even though it wasn't.

7 And this is still a recollection, so we'd
8 have to look at it much closer.

97

9 Q. Sure. Do you know when you got to the point
10 where you decided to do all the read multiple commands
11 with prefetches?

12 MR. TETER: Object to form.

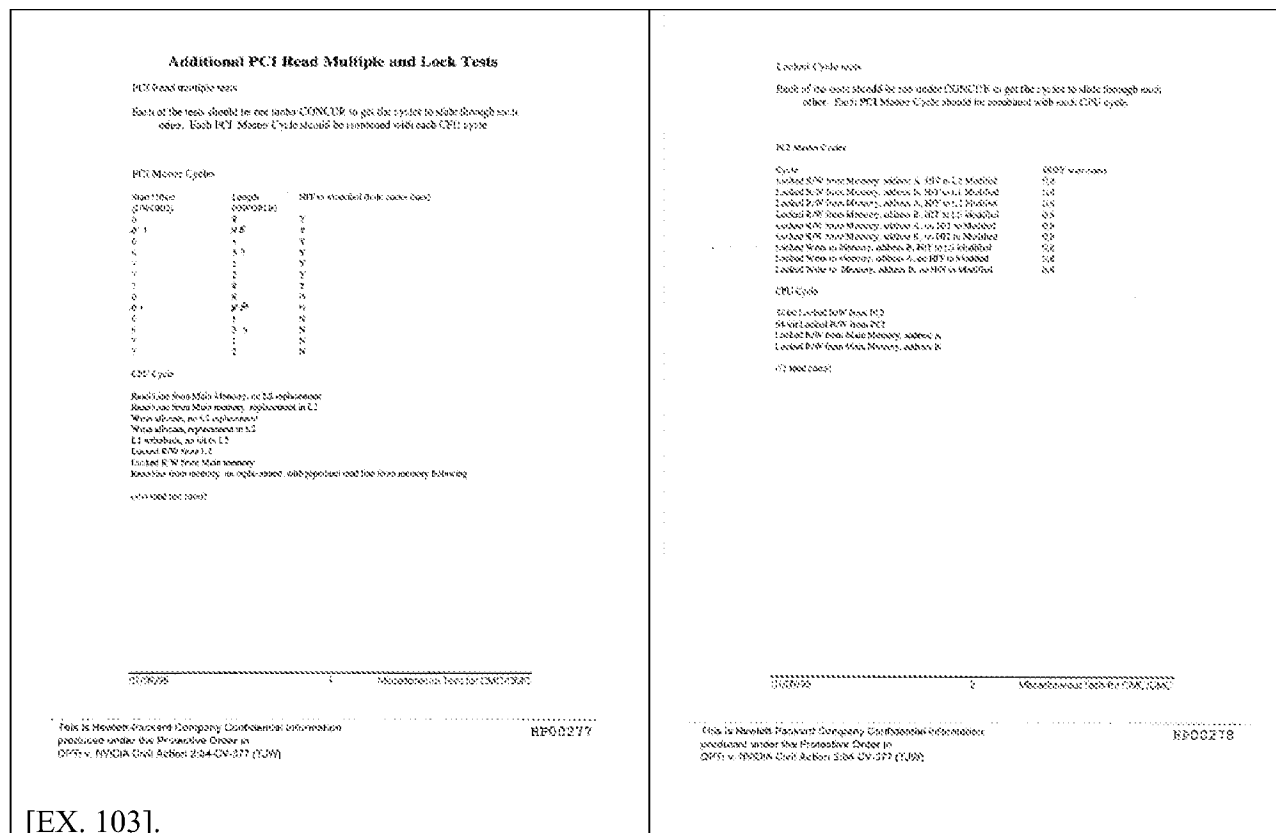
13 A. It would have been fairly early in the -- the
14 development cycle that it was pretty clear from the PCI
15 specification that the read multiple command was
16 intended to provide a hint to us that the PCI device was
17 intending to continue reading.

18 And so it made obvious sense to us that we
19 would, in the process of developing, want to go ahead
20 and -- and read ahead.

348. Even if the typewritten documents were read as OPTi counsel interprets them, they are not relevant under the jury instructions. Mr. Thome and Mr. Collins, the inventors, planned to burst across cache line boundaries from the outset. Mr. Thome testified that at most someone may have for a time proposed a different approach, but the inventors, Mr. Thome and Mr. Collins, conceived of bursting across cache line boundaries.

9. The documents that OPTi cites do not contradict Compaq's reduction to practice.

349. During the depositions of the Compaq inventors, OPTi's counsel argued that a Compaq document, entitled "Additional PCI Read Multiple and Lock Tests" shows that Compaq had not reduced the MRM method to practice even as late as January 1995, the date printed on the document. The document cited by OPTi appears below.



1 versions of the document, also.

2 Q. Okay.

3 A. Since the document -- documented -- we would
4 have documented the -- the tests that we were doing
5 probably over a series of time, but I don't recall the
6 particulars of it.

7 Q. And then Exhibit 103 would appear to indicate
8 that, as of January 1995, these -- these PCI Memory Read
9 Multiple tests were going to be run at some future date
10 under the CONCUR program, right?

11 MR. TETER: Object to form.

12 A. Or they were already previously run.

13 Q. (BY MR. TETER) Well, if you look at the -- the
14 first line of text, it says, "PCI" -- it says, "PCI Read
15 Multiple tests."

16 And then it says, "Each of the tests should
17 believe run under CONCUR." See that?

18 A. Yeah.

19 Q. Okay. So that's saying, in the future, they're
20 going to be run under CONCUR?

21 A. Well, as I -- I guess that's what I'm -- I'm
22 not saying this very well. But whenever this document
23 was written and whoever wrote it -- and I don't even
24 remember who wrote it -- and it may have even been me, I
25 don't remember -- when they wrote the document, they

1 would have written it in the future tense because they
2 were writing, saying these are the things that need to
3 be done. But as the document gets iterated and printed
4 out, the text probably wouldn't have changed but the
5 date would have been -- gone further and further out.

6 Q. Okay.

7 A. If that makes sense.

8 And I apologize. I just -- you know, like
9 I said, we don't seem to have the whole thing. If we
10 had the whole document, it would probably help me give a
11 little more context of what the particular situation
12 was.

22 Q. (BY MR. TETER) Does plaintiff's Exhibit 103
23 indicate to you that TriFlex/PCI in 1994 was unable to
24 run Memory Read Multiple?

25 A. I don't believe --

1 MR. BRODY: Object to the form.

2 A. I'm sorry.

3 MR. BRODY: I'm sorry. I just wanted to
4 object to form. Go ahead.

5 A. So I -- I don't believe that it -- it indicates
6 that PCI Memory Read Multiple did not run in 1994.

351. Based on all of the evidence cited above, OPTi's counsel appears to be misinterpreting the document. The document does not state that the MRM presnooping and prefetching feature did not work in February 1994.

10. Compaq conceived and reduced to practice over a lengthy period of time because it was making a very complex design, much more sophisticated than the relatively simple presnoop-only method described in the OPTi patent.

352. As discussed above, the OPTi presnooping method was much less sophisticated than the Compaq system or the Intel streaming method, and achieved results that fell far short of the Intel and Compaq chipsets. The OPTi presnoop patent, even as described in July 1995, failed to send data on every PCI clock, and as a result, was approximately half as fast as the Compaq CMC/DMC chipset and the Intel Triton. In view of the great complexity of the CMC/DMC, it is entirely consistent for Compaq to have conceived in late 1992, received silicon in February 1994, and released product in November 1994.

353. The deficiency in the OPTi design – its failure to send data on every clock cycle – suggests that the OPTi design was put together in a rush to compete with Intel and Compaq. The notion that OPTi would have conceived prior to Intel and Compaq, yet failed to figure out how to send data every clock cycle, is not plausible.

354. In October, 1994, Compaq had constructively reduced to practice an invention that included presnooping and prefetching, as set forth in Compaq's patent—Compaq filed the patent application that day, and I am informed that a patent application is a constructive reduction to practice of the methods and apparatus described therein.

355. Compaq/HP has provided clear and convincing evidence that it had reduced the presnooping feature in the CMC/DMC to practice by February 1994, when Compaq/HP had working silicon and was in the debugging process.

356. Compaq/HP has provided clear and convincing evidence that the conception of the Triflex/PCI prefetching/presnooping invention had occurred prior to September 1993, the date of the Biner chipset drawings, and prior to August 1993, the date of the feature comparison. [EX. 188].

357. The Compaq inventors indicated that they worked continuously on the Triflex/PCI until it was reduced to practice in silicon, and as well in the Compaq patent application.

358. OPTi has not provided any evidence demonstrating that it conceived of the presnooping invention before the Compaq inventors. Rather, OPTi appears to have had great difficulty implementing the invention, and as late as July 1995 had not yet managed to make a PCI chipset that could send a data unit on every clock cycle. As a result, the OPTi invention trailed the conception and reduction to practice of the Mr. Thome and Mr. Collins in Houston, Texas.

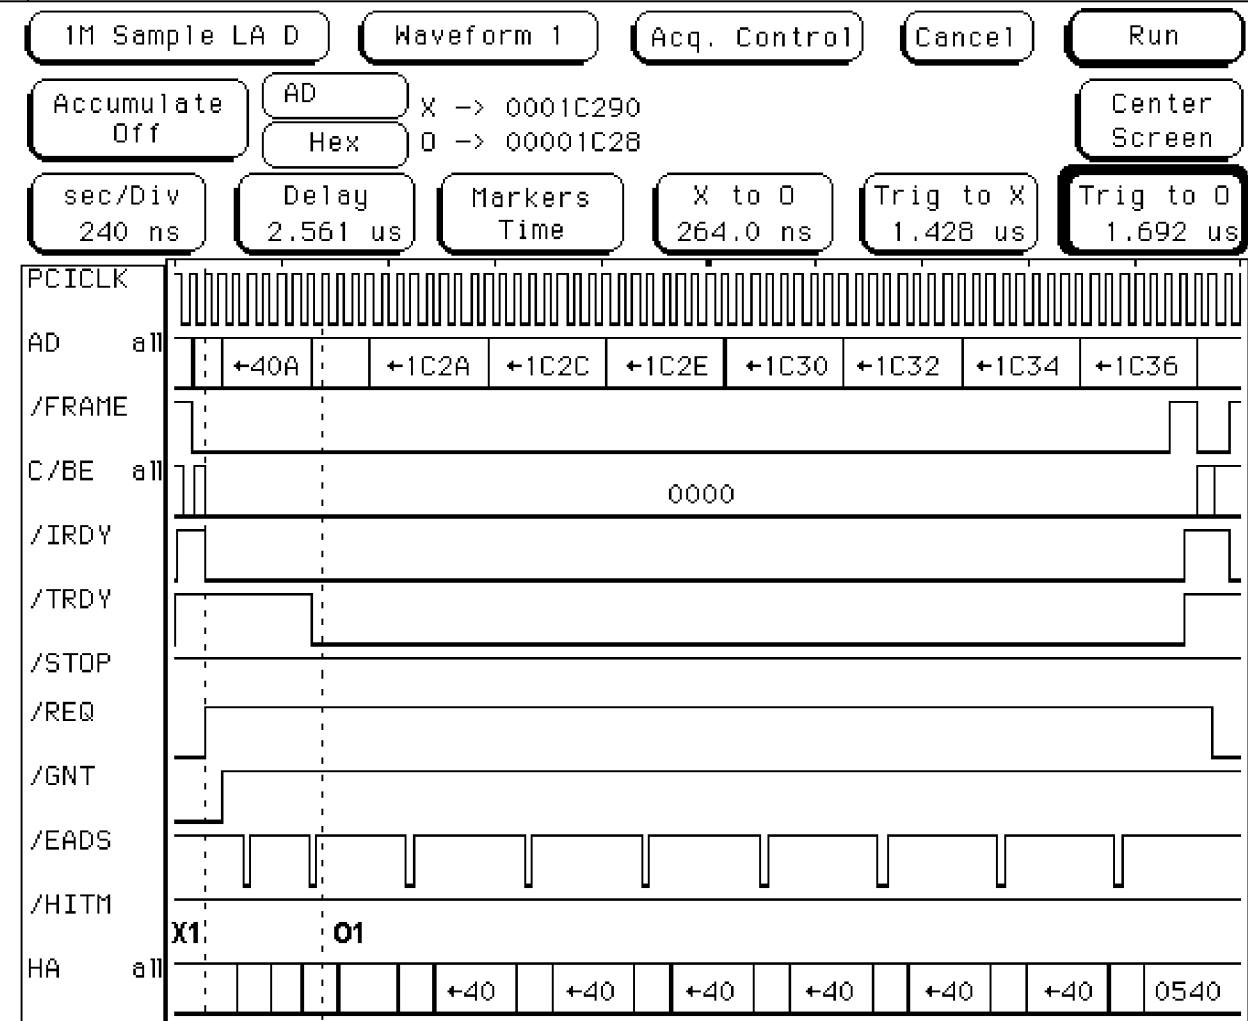
11. The Kramer Report shows that the Prosignia operated as the Compaq/HP witnesses testified.

359. The Compaq/HP evidence indicated that the Prosignia 500 included the presnooping and prefetching, allowing for a burst across cache line boundaries. In addition to evidence from Mr. Collins, Mr. Thome, and HP/Compaq itself, I have reviewed Silicon Valley Expert Witness Group (SVEWG's) Engineering report on the Compaq chipset ("Kramer report"). See Attachment C. The Kramer report shows that the Compaq chipset allowed streaming across cache line boundaries. Significantly, the Kramer report shows that the Compaq chipset implemented streaming across cache line boundaries by keeping the signals IRDY (initiator ready), TRDY (target ready), and FRAME asserted throughout the burst transaction. See Attachment C.

360. In addition, the Kramer report shows that the Prosignia 500 snooped only one line ahead, as recited in FIG. 9 of the '906 patent. Thus, the trigger for the next line

inquiry in the Prosignia 500 is clearly linked to the start of the transfer of the current line on the PCI bus. Figure C4 from the Kramer report, Attachment C, follows:

Figure C4. An example of a PCI master-initiated burst read transfer where the first inquiry cycle discovers the desired line is not cached cache modified in L1 cache.



The data being transferred over PCI (from the test buffer in main memory) set equal to its corresponding cache tag address (HA[17:5]) plus one 0 bit (corresponding to HA[4]) appended to the least significant end.

This clarifies the temporal relationship between the cache line tag addresses being snooped (with EADS# and HA) with the data in that cache lines that are being transferred over the PCI bus (AD). Note that the addresses being snooped (on HA at the time of the EADS# assertion) are not visible in this diagram because of its time scale. The values ending in "40" are *not* the addresses. Refer to Figures C4a-d to see the snoop address values displayed on the timing analysis.

(X1) is the address phase. This transfer is a PCI bus mastered Memory Read Multiple burst read from address 0x0001C290. (O1) is the first data phase (the first dword transfer).

[C4.bmp]

361. I should note that Mr. Kramer's analysis, though of a Prosignia 500, addresses a computer that had a Intel Pentium CPU at 150MHz, and that the BIOS had a date of July 1996. I understand that NVIDIA and Mr. Kramer have to date been unable to locate a Prosignia 500 with an earlier BIOS, and reserve the right to supplement this report in the event that Mr. Kramer and NVIDIA locate another Prosignia 500. In any event, I have no reason to believe that the presnooping and prefetching of the Prosignia 500 released in November 1994 would be any different than the presnooping and prefetching of the machine that Mr. Kramer tested and analyzed.

12. Acacia's analysis of the Compaq prior art is consistent with the rest of the evidence.

362. I have reviewed the transcript of Mr. Harris of Acacia. I understand that Acacia was considering buying the OPTi presnoop patents, and did not because Acacia identified the prior art Compaq patent. Mr. Harris testified about the discovery of the Compaq patent as follows:

31

20 Q And did you and Mr. Roop ever talk about this
21 Compaq patent?

22 A Yes, we did.

23 Q And what did you guys talk about?

24 MR. GRACEY: Object to form.

25 THE WITNESS: I asked Mr. Roop what his thoughts were

32

1 regarding potential invalidity.

2 BY MR. BRIGHAM:

3 Q And what were his thoughts?

4 MR. GRACEY: Object to form.

5 THE WITNESS: He thought that we would have a problem if
6 we tried to license the patents with regard to validity.

7 BY MR. BRIGHAM:

8 Q Because of the Compaq patent?

9 A Correct.

10 MR. GRACEY: Object to form.

11 BY MR. BRIGHAM:

12 Q And do you recall what you said?

13 MR. GRACEY: Object to form.

14 THE WITNESS: No, I don't recall.

15 BY MR. BRIGHAM:

16 Q Do you recall what your reaction was?
17 MR. GRACEY: Object to form.
18 THE WITNESS: Not right then.
19 BY MR. BRIGHAM:
20 Q You don't recall what your reaction was then?
21 A We decided, once we reviewed the patent, not to
22 make another offer.
23 Q So, did anything happened next with respect to
24 OPTi?
25 A I believe we pretty much stopped spending money
33
1 and time on that opportunity.

101
5 Q So, the only remaining factor was the validity
6 analysis; is that right?
7 A That was the issue that made us to decide not to
8 spend any more time with it.
9 Q Right.
10 But that was the issue that whether you went
11 forward with a higher offer or not hinged upon or whether you
12 went directly to the shareholders; is that right?
13 A For us, that was the showstopper.

363. Although I do not rely on the Acacia analysis per se, I agree with Acacia's conclusion. The Compaq prior invention is clear and convincing evidence that the asserted presnoop claims are invalid.

13. Compaq claim chart & timeline.

364. A claim chart is attached hereto as Attachment F, illustrating that Compaq's prior invention renders the presnoop patent claims-in-suit invalid. A timeline is attached hereto as Attachment G.

C. The asserted claims of the OPTi patents are invalid in view of the prior invention by Motorola inventors of the Motorola MPC105, and the MPC105 articles.

365. Between 1992 and 1994, Motorola conceived and designed a PowerPC-to-PCI bridge/memory controller chip called the Motorola MPC105. The MPC105 served as an interface between a PCI bus and system memory, and implemented both presnooping and prefetching features to facilitate a PCI bus master burst read of system memory. Motorola

conceived and implemented presnooping as described and claimed by the Patents-in-Suit by at least March of 1994.

- 1. The design of the MPC105 began in late 1992 or early 1993.**

2. The design of the MPC105 had been completed by at least March of 1994 when Motorola began submitting papers describing the MPC105 to academic and commercial conferences.

372. Mr. Bryant's testimony is consistent with the documentary evidence produced by Alan J. Smith. Alan J. Smith, who has been retained by OPTi in this case, was an organizer of the Hot Chips VI Symposium in August of 1994. [See SMITH000047-000049]. Motorola presented the MPC105 at the Hot Chips VI Symposium (see EX. 66, NV020000-020007), and the organizers of the Symposium had placed an advertisement in the February 1994 *Newsletter of the Parallel Processing Technical Committee* asking for the submission of papers by March 24, 1994.

Hot Chips VI, August 14-16, 1994, Stanford University, Stanford, California, USA. Send an extended abstract (1-3 pages) by March 24, 1994 to Prof. Alan Jay Smith, Computer Science Division, UC Berkeley, Berkeley, CA 94720, Fax: (510)-539-6360, E-mail: smith@cs.berkeley.edu. Submissions can be made by hard copy, fax, or electronic mail (E-mail is preferred!).

[EX. 67 at NV 019294 (Newsletter of the Parallel Processing Technical Committee at 12)].

373. Mr. Bryant recalled that a presentation about the MPC105 was made at "Hot Chips" and believed that the MPC105 had "taped out" by the deadline for submission of papers:

31

16 Q. Were you ever the author of any articles
17 describing MPC105?

18 A. I helped with authoring articles, yes.

19 Q. Do you know if any of the articles were ever
20 present at any conferences?
21 A. Yes.
22 Q. Do you know whether there were slides associated
23 with presentations at conferences?
24 A. Uh-huh. Hot Chips.

35

3 Q. Do you have any reason to believe that an
4 extended abstract was not sent to Professor Alan Jay
5 Smith by March 24, 1994?
6 A. No.
7 Q. Do you believe that the MPC105 had taped-out by
8 March 24, 1994?
9 A. Yeah.

374. Alan J. Smith produced documents which confirm that another designer on the MPC105, Karl Wang, did indeed submit a paper describing the presnoop and prefetch features of the MPC105 by March 24, 1994:

I am sending you by this e-mail an abstract of our paper we like to submit to the Hot Chips Symposium VI. This abstract contains slight modification from the abstract which I send to you by fax on March 24. Please accept this modified version. My appology

[EX. 81 at SMITH000023 (3/25/94 Email from Karl Wang to Alan J. Smith)].

375. The paper that Karl Wang submitted to Alan J. Smith contained a description of the presnooping and prefetching features implemented in the MPC105:

forwarded to memory. Finally, EAGLE implements a 32-byte PCI-to-memory read prefetch buffer. If speculative reads are enabled, at the completion of a PCI read EAGLE prefetchs the next cache line from memory to enhance the PCI-to-memory read performance. Since the processor, secondary

[EX. 81 at SMITH000024 (3/25/94 Email from Karl Wang to Alan J. Smith)].

376. Moreover, the mail from Karl Wang to Alan J. Smith specified the status of the project as "first silicon":

Status of Product: "first silicon" Product not yet announced. Please keep submission confidential.

[EX. 81 at SMITH000028 (3/25/94 Email from Karl Wang to Alan J. Smith)].

377. Mr. Bryant confirmed that this sequence of events was consistent with his recollection:

36

- 3 Q. I would like to mark as Exhibit 81 a document
4 labeled Smith 000023 to 000028.
5 (Defense Exhibit No. 81 marked.)
6 Q. Mr. Bryant, have you seen this document before?
7 A. No.
8 Q. Do you know who Karl Wang is?
9 A. Yes.
10 Q. Who is Karl Wang?
11 A. Karl Wang was the Central Control Unit team lead
12 of the MPC105.
13 Q. What were his responsibilities?
14 A. Overseeing the core of the chip design.
15 Q. Do you remember him sending this e-mail for
16 referencing its sending to Professor Alan J. Smith in
17 March of 1994?
18 A. No.
19 Q. Do you remember contributing to an abstract that
20 was sent to Professor Alan J. Smith?
21 A. I remember Karl asking me questions about how PCI
22 worked. I didn't know what he wanted it for.
23 Q. Did you write any part of the abstract?
24 A. Not that I recall.
25 Q. If you look at Page 2 of Exhibit 81, this is a

37

- 1 document labeled Smith 000024. Down near the bottom,
2 second to last sentence reads, "If speculative reads are
3 enabled at the completion of the PCI read EAGLE
4 prefetches the next cache line from memory to enhance
5 the PCI to memory read performance." Do you see where I
6 am referring to?
7 A. Yes, I do.
8 Q. What does "EAGLE" refer to?
9 A. EAGLE refers to our code name for the MPC105.
10 Q. And the time this e-mail was written on or around
11 March 26th of 1994, was that an accurate description of
12 how the MPC105 code named EAGLE behaved?
13 A. Yes.

3. **Motorola received a working silicon chip for the MPC105 on May 22, 1994.**

4. After reducing the MPC105 to silicon, Motorola continued to publish articles and prepared to sell the MPC105 commercially.

379. Motorola also participated in and presented the MPC105 at the ICCD 94 International Conference on Computer Design in October of 1994. [See EX. 61 at OPTINVIDIA 000681 (IEEE International Conference on Computer Design: VLSI in Computer and Processors, *Proceedings* at 412)]. The documentary evidence shows that the deadline for submission of papers to this conference was March 4, 1994, with a final manuscript due by July 1, 1994.

CALL FOR PAPERS

**ICCD 94 International Conference on Computer Design,
October 10-12, 1994, Royal Sonesta Hotel, Cambridge,
Massachusetts. Original papers regarding work in the fol-
lowing areas are being solicited:**

- Embedded systems
- VLSI & technology
- Architecture and algorithms
- Design & test
- Computer-aided design

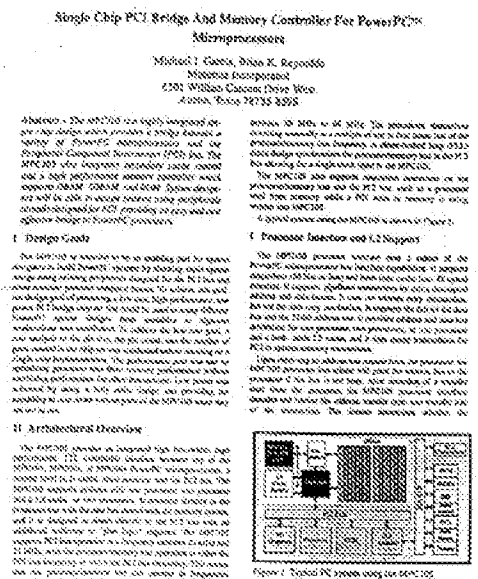
**Submissions should be sent to technical program chair
A.P. Ambler, Dept. Electrical Engineering & Electron-
ics, Brunel University, Uxbridge, Middlesex. UB8 3PH,
+44 895-203380, fax +44 895-258728, tony.ambler@
brunel.ac.uk. Deadline for submissions is March 4, 1994,
with notification of acceptance by April 30 and final man-
uscript due July 1.**

[EX. 62 at NV 019301 (Spring 1994 *Design & Test of Computers* at 2)].

380. The paper Motorola presented at the conference (see EX. 61 at OPTINVIDIA 000676-000681) also clearly described the MPC105's presnooping and prefetching features.

loaded. The MPC105 also employs a speculative read feature which, if enabled, will snoop the next sequential cache line address when the current PCI read is accessing the third doubleword (8 bytes) of the cache line. Once the speculative snoop response is known and the read has finished, the data at the speculative address will be fetched and loaded into the buffer in anticipation of the same PCI master requesting it. If it is not requested, the data is just invalidated. In a typical system, the speculative feature will allow streaming of large amounts of data with very minimal delay on the PCI bus enabling the system to utilize the full bandwidth of the PCI bus.

[Ex. 61 at OPTINVIDIA 000681 (Proceedings, IEEE International Conference on Computer Design: VLSI in Computer and Processors at 412)].



381. The testimony of Mr. Bryant and Mr. Garcia, along with the calls for papers and the documentary evidence produced by Alan J. Smith, prove that the MPC105 design was completed and “taped out” in March 1994 at least, and that the design team had begun submitting papers describing the features of the MPC105 at that time. This evidence establishes that the MPC105 design team had conceived of presnooping and prefetching by at least March 24, 1994.

5. Motorola presented the MPC105 at Hot Chips VI in mid-August 1994 and formerly announced the MPC105 on August 22, 1994.

382. As Mr. Bryant testified, Motorola presented the MPC105 at the Hot Chips VI Symposium at Stanford University on August 14-16, 1994. [See SMITH000047-000269]. The testimony of Brian Reynolds confirms Mr. Bryant’s recollection that the MPC105 was presented at Hot Chips:

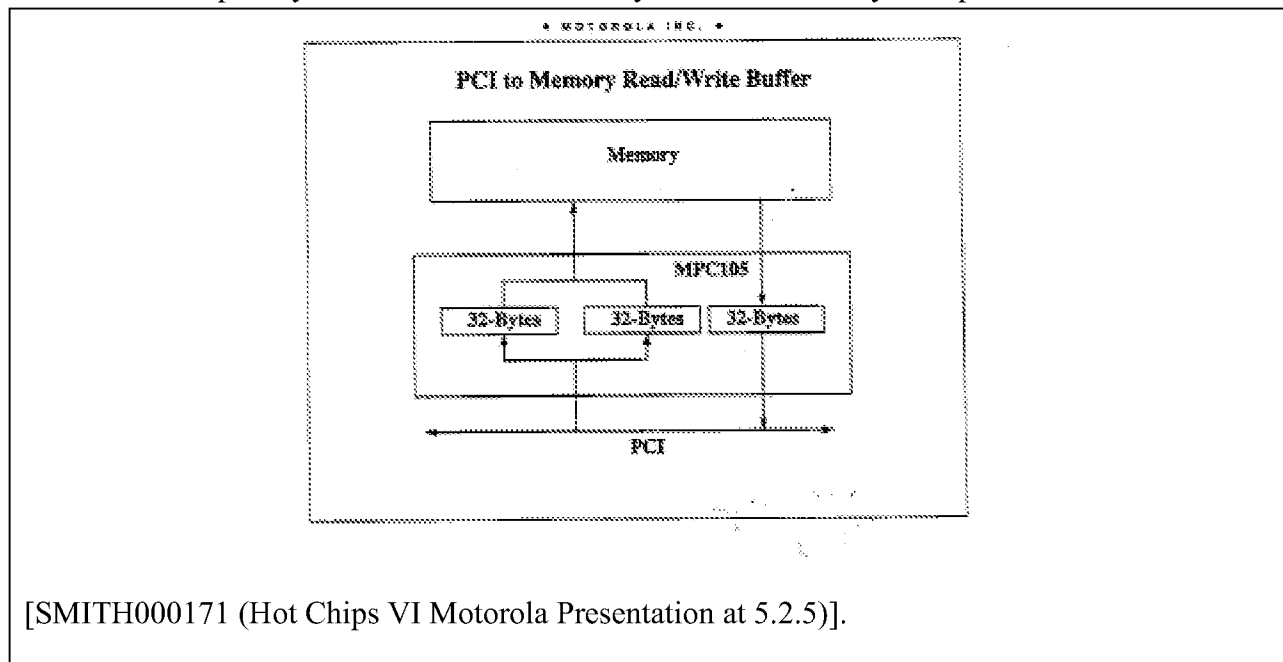
- 9
- 10 Q. Have you heard of the Hot Chips VI symposium
- 11 at Stanford?
- 12 A. Yes.
- 13 Q. And what was the Hot Chips VI symposium at
- 14 Stanford?

15 A. It's a technical design conference that people
 16 throughout the chip/semiconductor industry would attend
 17 and present some of their latest works.
 18 Q. Did you present anything at the Hot Chips VI
 19 symposium at Stanford University in August of 1994?
 20 A. No.
 21 Q. Did any of your colleagues?
 22 A. I believe Karl Wang presented. He was one of
 23 the design managers on the project.
 24 Q. Do you know what Karl Wang presented in August
 25 '94 at the Hot Chips VI symposium at Stanford?

10

1 A. Not for sure, but I believe it was a
 2 presentation based on a paper that had been written
 3 about the MPC105.

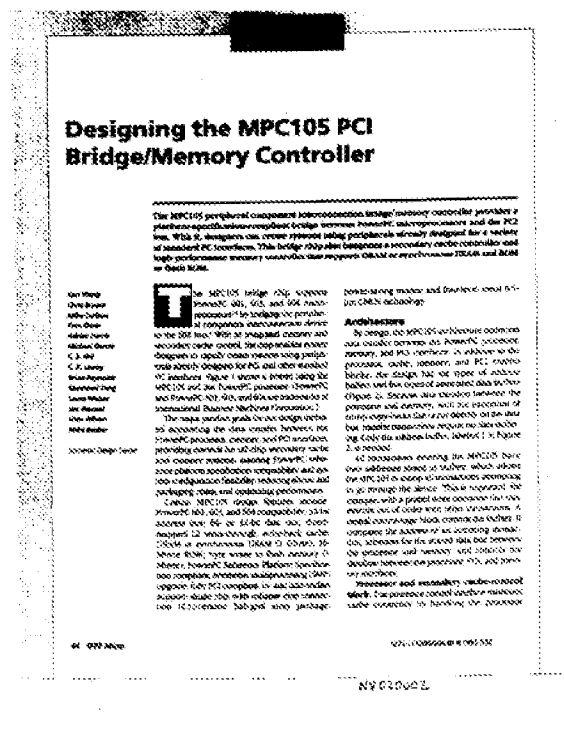
383. In addition, Alan J. Smith has produced Motorola's presentation at the Hot Chips VI Symposium. [See SMITH000167-000174]. The presentation details the MPC105 features and explicitly shows the MPC105 32-byte PCI-to-Memory read prefetch buffer:



384. The presentation that Motorola inventor Karl Wang presented at the Hot Chips IV Symposium, and the subsequent paper he published regarding it, describes the MPC105 presnoop and prefetch feature in accordance with the testimony and documentary evidence that I have cited in this report:

The MPC105 also implements a speculative read feature that, if enabled, will snoop the next sequential cache line address when the current PCI read accesses the third doubleword of the cache line. Once the snoop response is known and the read has finished, the device will fetch the data at the speculative address and load it into the buffer in anticipation of the same PCI master requesting it. In a typical system, the speculative feature will allow streaming of large amounts of data with very minimal delay on the PCI bus, enabling the system to use the full bandwidth of the PCI bus.

[EX. 66 at NV020005 (Hot Chips IV Symposium Paper at 47)].



385. In the summer of 1994, Motorola announced the MPC105 PowerPC-to-PCI bridge/memory controller. An article published by Electronic News on August 22, 1994 stated that the Motorola MPC105, code-named "Eagle", would be incorporated into a Motorola motherboard supporting PowerPC 603 and 604 processors. The announcement specifically referenced the presentation of the MPC105 at the Hot Chips VI Symposium, and stated that evaluation boards were then available, with manufacturing kits to be available in October of 1994:

[EndArticles](#) > [Electronic News](#) > [August 22, 1994](#) > [Article](#) > [Print Friendly](#)

Motorola PowerPC move signals rivalry with IBM - Motorola rolls out reference motherboard, MPC105 bridge/memory controller - Product Announcement

Reinhardt Krause

AUSTIN, TEX.--Marking the point where it will seriously begin competing for customers with IBM, Motorola's RISC Microprocessor division last week began offering its own PowerPC reference motherboard and developer's program--as well as a new single-chip PowerPC-to-PCI bridge/memory controller.

While Motorola's reference motherboard is compliant with IBM's reference platform specification there are some differentiators. One is the new PowerPC-to-PCI bridge/memory controller--the MPC105 code-named Eagle--which was described at the Hot Chips VI symposium at Stanford University in Stanford, Calif. last week (see related stories, page 2). While immediate comparisons may be with Intel's Peripheral Component Interconnect (PCI) chipsets, the new Eagle bridge/memory controller also will compete with IBM's two-chip implementation, the 82650, unveiled last year (EN, Nov. 1, 1993).

...
...

An evaluation kit is available now from Motorola for \$25; a manufacturing kit will be shipping in October for \$50. Motorola will initially be providing ARC (Advanced RISC Compliant) firmware and HAL (hardware abstraction layer) support and not the open-boot capability that IBM is still developing.

[EX. 59 at NV 012088-012089].

6. **Motorola made additional presentations of the MPC105 in the latter half of 1994 and expected to begin offering the MPC105 for sale at the end of 1994.**

386. As I cited above, Motorola again presented the MPC105 at the IEEE International Conference on Computer Design on October 10-12, 1994 and once again described the presnoop and prefetch features:

loaded. The MPC105 also employs a speculative read feature which, if enabled, will snoop the next sequential cache line address when the current PCI read is accessing the third doubleword (8 bytes) of the cache line. Once the speculative snoop response is known and the read has finished, the data at the speculative address will be fetched and loaded into the buffer in anticipation of the same PCI master requesting it. If it is not requested, the data is just invalidated. In a typical system, the speculative feature will allow streaming of large amounts of data with very minimal delay on the PCI bus enabling the system to utilize the full bandwidth of the PCI bus.

[EX. 61 at OPTINVIDIA000676; OPTINVIDIA000681 (IEEE International Conference on Computer Design: VLSI in Computers and Processors, *Proceedings at 412*)].

Single Chip PCI Bridge And Memory Controller For PowerPC™ Microprocessors

Michael J. Davis, Brian C. Reynolds
Motorola Incorporated
4001 Williams Court Drive West
Austin, Texas 78758-0538

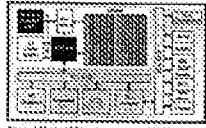
Abstract - The MPC105 is a single-chip integrated circuit which provides a bridge between a PowerPC™ microprocessor and the PCI bus. The MPC105 also provides a memory controller and a 32-bit performance monitor. The MPC105 is a 3.3V CMOS device which is available in a 100-pin Quad Flat Pack (QFP) package.

1. Bridge Description

The MPC105 is a single-chip integrated circuit which provides a bridge between a PowerPC™ microprocessor and the PCI bus. The MPC105 also provides a memory controller and a 32-bit performance monitor. The MPC105 is a 3.3V CMOS device which is available in a 100-pin Quad Flat Pack (QFP) package.

2. Performance Monitor

The MPC105 includes a 32-bit performance monitor which can be used to monitor the performance of the microprocessor. The monitor can be configured to monitor a variety of performance metrics, including the number of instructions executed, the number of cache hits, and the number of bus transactions.



OPTINVIDIA 000676

387. Mr. Reynolds testified that Motorola customers had operating boards prior to October 12, 1994:

- 149
- 24 Q. Okay. You said you thought that customers had
- 25 operating boards at some point prior to October of
- 150
- 1 1994 -- or, I guess, October 21st of 1994.
- 2 A. I believe so, yes.
- 3 Q. Do you know how far in advance?
- 4 A. No.
- 5 Q. You said that the boards that were shipped
- 6 were capable of performing the speculative read and
- 7 speculative snoop functions as described. Did I hear
- 8 that correctly?
- 9 A. Yes.

388. An article in Microprocessor Report on September 12, 1994, reported that Motorola expected shipments of the MPC105 to begin at the end of 1994:



Newspapers

Aiming to simplify the design of PowerPC systems, Motorola and its partners have fired off a flurry of announcements. The most significant is that Motorola will market a system-logic chip, the MPC105, that will provide cache control, memory control, and a PCI interface for any PowerPC 60x processor. Although this chip is the second PowerPC system-logic design to be announced, it is the first with such a high degree of integration. Motorola expects to begin shipments by the end of this year.

[EX. 71 at NV20008].

7. The Motorola MPC105 was designed to operate as an interface between PowerPC 60x processors and a PCI bus.

389. Generally speaking, the Motorola MPC105 was a microchip interface between a PowerPC 60x CPU, a secondary (L2) cache, the PCI bus, and main memory. [See EX. 65 at NV 012239 (MPC105 User Manual at 1-1)]. The MPC105 connected to and communicated with each of these components and facilitated data transfers between them. The MPC105 is placed in context in Exhibit 63, which is an enlarged copy of Figure 1 from Motorola's publication at the IEEE Conference on Computer Design:

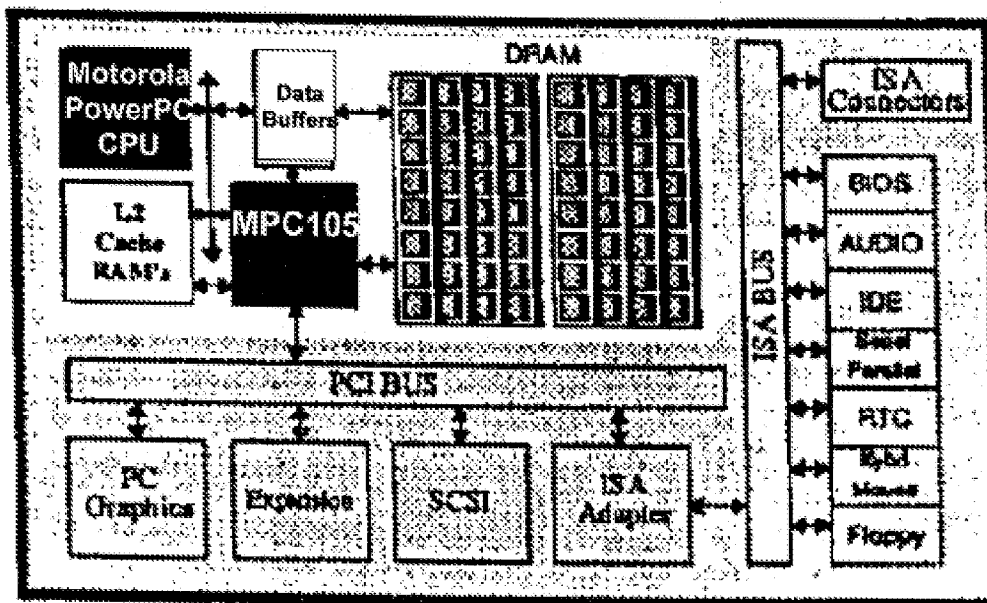
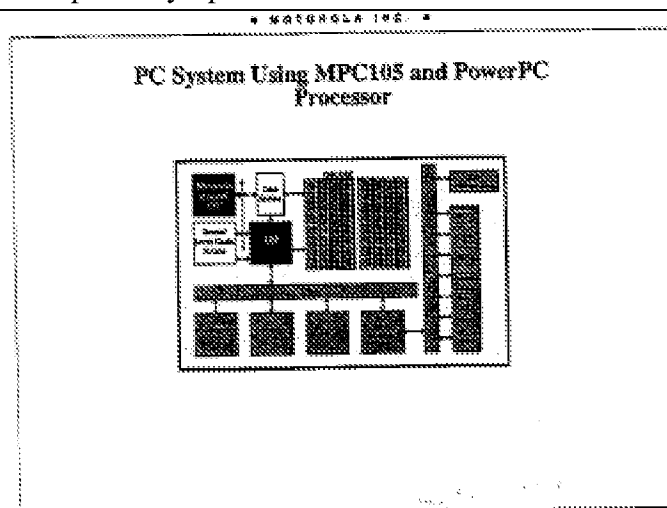


Figure 1 Typical PC system using the MPC105.



[EX. 63, (enlarged figure from Ex. 61, OPTINVIDIA 000678)].

390. The figure above was also displayed in Motorola's August 1994 presentation at the Hot Chips VI Symposium at Standard:



[SMITH00168 ((Hot Chips VI Motorola Presentation at 5.2.2)).

8. The MPC105 relied upon and implemented portions of the PCI Local Bus Specification, Revision 2.0, and also implemented presnooping as described in the Patents-in-Suit.

391. One of the primary purposes of the Motorola MPC105 was to facilitate data transfers between the PCI bus and main memory. [See EX. 65 at NV 012439-012459 (MPC105 User Manual, Ch. 7)]. In particular, the MPC105 user manual states that the MPC105 responded to PCI bus commands, including the PCI Memory Read Multiple command:

Table 7-1. PCI Bus Commands (Continued)				
C/BE3- C/BE0	PCI Bus Command	MPC105 Supports as a Master Device	MPC105 Supports as a Target Device	Definition
1011	Configuration-write	Yes	No	The configuration-write command accesses the 256-byte configuration space of a PCI agent. A specific agent is selected when its IDSEL signal is asserted during the address phase. See Section 7.4.5.2, "Accessing the PCI Configuration Space," for more detail of configuration accesses.
1100	Memory-read-multiple	No	Yes	The memory-read-multiple command functions the same as the memory-read command on the MPC105. If prefetching is desired, speculative PCI reads should be enabled. See Section 8.1.3.2.1, "Speculative PCI Reads from System Memory," for more information.

[EX. 65 at NV 012444 (MPC105 User Manual at 7-6)].

392. As I have stated previously in this report, the PCI Memory Read Multiple command was designed as a hint that the bus master wishes to read more than one cache line in a transaction. The PCI Local Bus Specification, Revision 2.0 describes the Memory Read Multiple command as semantically identical to the PCI Memory Read command "except that it additionally indicates that the master may intend to fetch more than one cache line before disconnecting." In addition, the description of the PCI Memory Read Multiple command

suggests that the memory system “might gain some performance advantage by sequentially reading ahead an additional cache line . . .”.

The *Memory Read Multiple* command is semantically identical to the *Memory Read* command except that it additionally indicates that the master may intend to fetch more than one cache line before disconnecting. The memory controller should continue pipelining memory requests as long as *FRAME#* is asserted. This command is intended to be used with bulk sequential data transfers where the memory system (and the requesting master) might gain some performance advantage by sequentially reading ahead an additional cache line when a software transparent buffer is available for temporary storage.

[EX. 64 at OPTINVIDIA 001546 (PCI Local Bus Specification, Revision 2.0 at 21)].

393. Indeed, the MPC105 responded to the PCI Memory Read Multiple Command by presnooping and prefetching. The MPC105 User Manual describes the presnoop as a “speculative snoop” and the prefetch as a “speculative read”: “When this [speculative read] feature is enabled . . . the MPC105 starts the snoop of the next sequential cache line address when the current PCI read is accessing the third double word (the second half) of the cache line in the PCMRB [PCI-Read-from-System-Memory-Buffer].” [Ex. 65 at NV 012468 (MPC105 User Manual at 8-8)].

8.1.3.2.1 Speculative PCI Reads from System Memory

To minimize the latency for large block transfers, the MPC105 includes a selectable speculative read feature. When this feature is enabled ($PICR1[2] = 1$), the MPC105 starts the snoop of the next sequential cache line address when the current PCI read is accessing the third double word (the second half) of the cache line in the PCMRB.

Once the speculative snoop response is known and PCI has completed the read, the data at the speculative address is fetched from system memory and loaded into the buffer in anticipation of the next PCI request. Note that the assertion of CAS is delayed until PCI is finished reading the data currently latched in the buffer. If a different address is requested, the speculative operation is halted and any data latched in the PCI read buffer is invalidated.

[EX. 65 at NV012468 (MPC105 User Manual at 8-8)].

394. Each of the three different MPC105 designers that were deposed testified that the descriptions of presnooping and prefetching in the Hot Chips VI publication, the

IEEE publication, and the MPC105 User Manual were accurate and that the MPC105 operated as described. Michael Garcia testified that the MPC105 would snoop the next line while the current line was being transferred:

8

6 Q. Are you one of the designers of the Motorola MPC105
7 bridge chip?

8 A. Yes.

9 Q. Did the MPC105 bridge chip have a speculative snoop
10 feature?

11 A. Yes.

12 Q. Did the MPC105 bridge chip have a speculative read
13 feature?

14 A. Yes.

57

3 Q. And that system, while data was being transferred
4 from a current line, would speculatively snoop the next line; is
5 that right?

6 MR. BRODY: Object to the form.

7 THE WITNESS: Yes, if -- if it was to snoopable
8 memory space.

9 Q. (BY MR. TETER) What is snoopable memory space?

10 A. PCI has two different types of memory space, some of
11 it is snoopable and some of it is not, and you have to keep
12 coherency in the part of the memory that's snoopable. And the
13 part that's not snoopable, you don't care whether its coherent
14 or not.

15 Q. So to follow that up, the system with the MPC
16 described in Exhibit 61 would snoop snoopable memory for the
17 next line while data units from the current line were being
18 transferred to the PCI master; is that right?

19 A. Yes.

395. Brian Reynolds also testified that the MPC105 operated to presnoop and prefetch as described in the MPC105 User Manual:

7

3 Q. Are you one of the designers of the Motorola
4 MPC105 bridge chip?

5 A. Yes.

6 Q. And at a high level did the MPC105 bridge chip
7 have a speculative snoop feature?

8 A. Yes.

9 Q. Did the MPC105 bridge chip have a speculative
10 read feature?

11 A. Yes.

72

21 Q. Looking at Exhibit 68, and the two papers that
22 you published regarding the MPC105, did the MPC105
23 implemented in a computer system include a secondary
24 memory?

25 A. Yes.

73

1 Q. A cache memory?

2 A. Yes.

3 Q. PCI master?

4 A. Yes.

5 Q. A PCI master that could sequentially read data
6 units from secondary memory?

7 A. Yes.

8 Q. A snoop of the next line prior to the
9 completion of the transfer of data units from the
10 current line, did it include that?

11 A. Yes.

396. Christopher Bryant testified that he and the other designers implemented the Memory Read Multiple command in reliance upon its description in the PCI Local Bus Specification (EX. 64, Bryant Tr. at 16:3-19), and that the MPC105 consequently presnooped and prefetched in response:

21

14 Q. Did reading the PCI local bus specification
15 motivate the MPC105 team to implement a presnoop or
16 speculative read feature?

17 A. Yeah.

18 Q. What exactly motivated that implementation?

19 A. It is in the spec. You want to be PCI compliant
20 then you're going to follow what's in the spec.

19

4 Q. What is the difference between the PCI memory
5 read and the PCI memory read multiple commands?

6 A. There are none. The only difference is that it
7 is a hint of what the bus master intends to do.

8 Q. And what is that hint?

9 A. That it is going to want more than one cache line
10 of data.

11 Q. Did the MPC105 behave differently depending on
12 whether it received a PCI memory read or a PCI memory
13 read multiple command?

14 A. Yes. A memory read only fetched one cache of
15 line of data. The memory read multiple, we ended up
16 doing the speculative snoops and the speculative reads
17 either from the L1 cache or the memory control.

9. The MPC105 had only one prefetch buffer, so it had to disconnect at cache line boundaries.

397. The MPC105 had two write prefetch buffers but only one read prefetch buffer. Mr. Garcia testified that Motorola considered the buffer size, but decided that increasing buffer size didn't add much performance:

18

22 Q. My question may have been bad. Let me just make sure
23 I was right. So the MPC105 had two buffers for writes --

24 A. Correct.

25 Q. -- and then one buffer for reads?

19

1 A. Right.

2 Q. Why did the MPC105 have only one buffer for reads and
3 yet two buffers for write? Can you explain that?

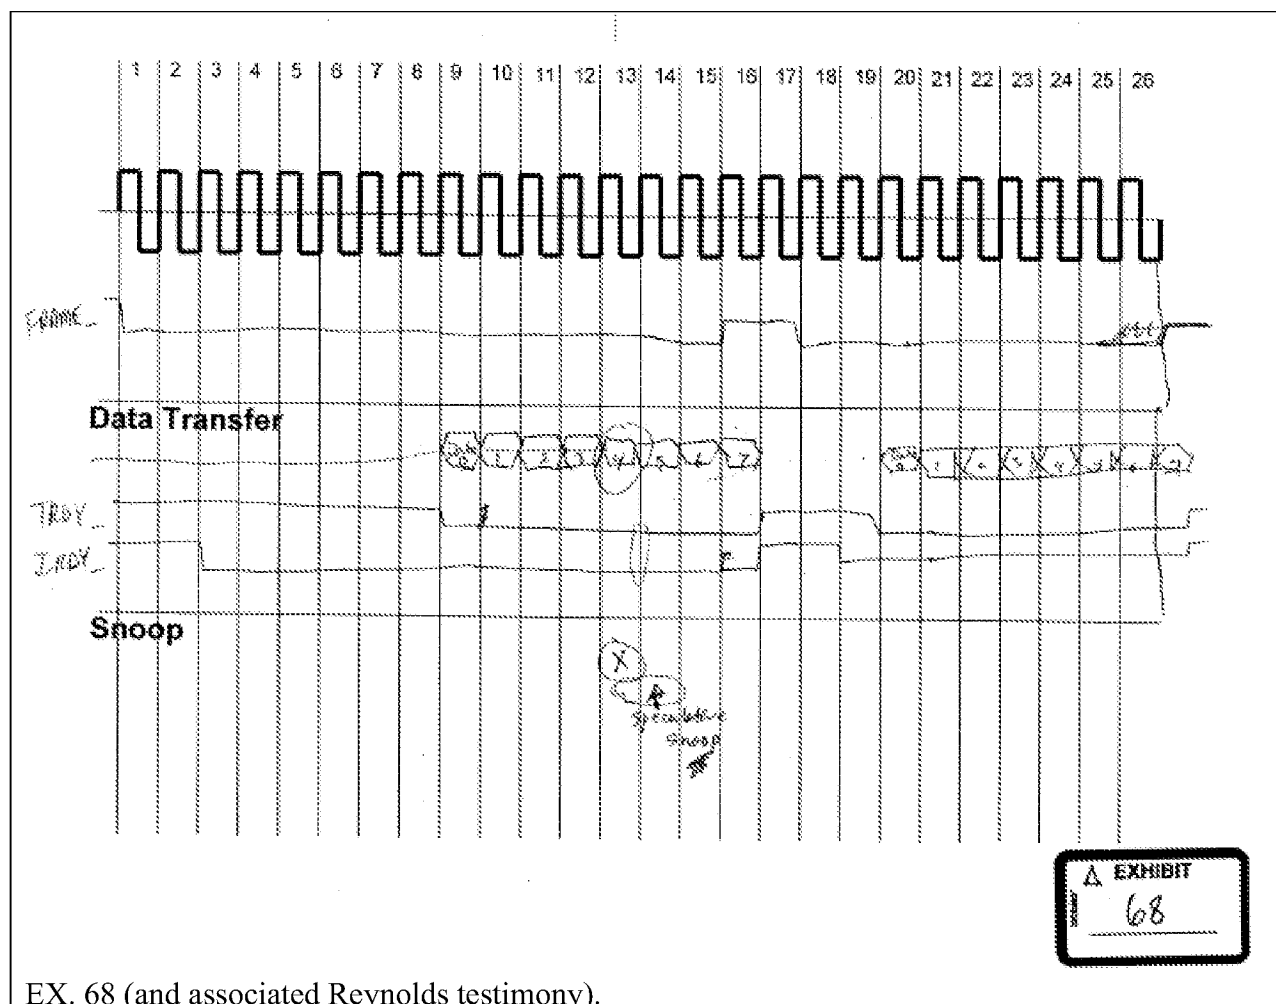
4 A. It probably had to do with performance and trying to
5 minimize the size of the chip and complexity. It could also
6 mean that we didn't need it for performance-wise. It didn't buy
7 us much to add an additional one.

19

9 The passage we just looked at indicates that
10 with the two buffers for writes, one of the buffers could be
11 being filled while the other buffer is forwarding to system
12 memory. Is that how it worked?

13 A. Yes.

399. Thus, the MPC would disconnect at cache line boundaries. At his deposition Mr. Reynolds drew a timing diagram showing a next-line inquiry, the data transfers, and the gap (or disconnect) between the transfer of the data units of the first line and the data units of the next line.



EX. 68 (and associated Reynolds testimony).

400. Although the MPC105 disconnected at cache line boundaries, as Mr. Reynolds testified, MPC believed that the performance was sufficient and adding an additional prefetch buffer would not be necessary or worth the extra cost:

140

11 Q. When you were working on the MPC105, did you
 12 understand that the reason there would be the
 13 disconnect at a cache line was because you only had a
 14 buffer that was one cache line long?

15 A. Yes. You might say that it was also
 16 because -- yeah. We didn't -- and also because we
 17 wouldn't partially fill it with the next line was the
 18 other reason we had to target disconnect.

19 Q. If you understood that the target disconnect
 20 was because you had a buffer that was one cache line
 21 long, why, then, didn't you just have two PCMRB
 22 buffers?

23 A. I think it was because the additional area and
24 complexity of controlling the second buffer was more
25 than we wanted to do. We also believed that the

141

1 performance benefit of the speculative reads was
2 sufficient for the trade-off of -- of the area for a
3 second one. It starts to -- I guess even with one it
4 looks like a cache. The more you have, the more
5 states -- or more you have to -- you have to track for
6 each individual one that stayed in the system. I was
7 going to say something else.

38

12 Q. Turning to the page prior, 411, in your
13 paper --

14 A. Uh-huh.

15 Q. -- the last paragraph, we looked at it before,
16 states, "There are three PCI" --

17 A. Right.

18 Q. -- "to system memory transaction buffers, one
19 for reads and two for writes." Do you see that?

20 A. Yes.

21 Q. Is that the way the MPC105 was actually
22 implemented?

23 A. Yes.

24 Q. Why did you design the system to have two
25 buffers for PCI writes to system memory, and one buffer

39

1 for PCI reads from system memory?

2 A. We had two buffers for PCI writes to system
3 memory so that we could post or store data that's being
4 written by PCI and not have to flush that buffer to the
5 main memory immediately if there is a subsequent write
6 on the PCI bus to use it. So the main goal was to
7 allow the maximum amount of data bandwidth on the PCI
8 bus.

9 And with the writes we determined that
10 we -- by the time the first write had written a full
11 cache line, we weren't really guaranteed that we could
12 have gotten the snoop done for that write and the data
13 pushed out of the buffer and into main memory. So we
14 had a second buffer to take up the slack, so while the
15 first one is doing its snoop and pushing data to
16 memory, the second one is being filled. And then it
17 would switch while the second one is being filled --
18 or, sorry, is doing its snoop and writing to memory the
19 first one would be filled.

20 On the read side, the -- because we were
21 doing the speculative snooping, we only felt we needed
22 one buffer -- the buffers take up area and add more
23 complexity to the chip, more possibility of errors.

24 But, anyway, the read one we determined -- had the same
25 goal, was to optimize the amount of bandwidth that you

40

1 could use on the PCI bus for data. And so we -- I
2 don't know why we didn't particularly use a second
3 buffer. We probably could have.
4 But it starts to look like another cache,
5 basically, because we were more storing more and more
6 lines of data as the owner. And so -- but mostly
7 was -- we decided the performance was sufficient to
8 only use the one buffer, use the speculative snooping
9 and prefetch feature, to get the second line of data in
10 without a significant loss in bandwidth on the PCI bus.

401. In addition, Mr. Reynolds testified, Motorola believed that there simply weren't bus master devices at the time that would use the MRM command.

144

3 Q. So was it -- was it your view at the time that
4 the speculative read feature allowed transactions to be
5 fast enough?

6 A. Yes. And at the time we also were not
7 convinced for the 105 that there were devices that were
8 using memory read multiple as a command. There
9 probably were, but it was not -- I don't think it was
10 prevalent. And so that also factored into our decision
11 to --

12 Q. Are you suggesting that at the time there
13 wasn't as much of a need for a bridge chip that would
14 allow for memory read multiple commands without any
15 kind of latency at all between data transfers?

16 A. Yes. Yes.

402. The motivation to modify the MPC105 to include a buffer that could avoid the disconnect could have come from any number of sources, including the PCI Design Guide EX. 97, which suggested prefetching (see discussion above in section re: Intel) or from the Amini patent, which taught the "ping-pong" structure of prefetch buffers. Using a "ping-pong" buffer structure, as taught by Amini, would have allowed the MPC105 to avoid any disconnect. In view of the cost/benefit considerations, the MPC105 inventors simply concluded that one buffer was enough for their purposes.

10. The MPC105's disconnect at cache line boundaries is not relevant to several of the '906 claims, and even under OPTi's interpretation, '906 claims 9 and 26 would be obvious in view of the MPC105 system.

403. '906 patent claims 1, 7, and 8 do not exclude a sequential transfer of data units where the PCI master disconnects at cache line boundaries, as in the MPC105. '906 patent claim 9 states that the transfers must be "at a constant rate," which OPTi apparently states would prohibit the PCI master from disconnecting at a cache line boundary, or otherwise preclude additional wait states between cache lines. "Constant rate" does not preclude wait states, as OPTi's own '906 patent does not transfer data on every PCI clock cycle. In any event, assuming that OPTi's construction is correct, and that no additional wait states can appear at a cache line boundary, claims 9 and 26 would be obvious – Motorola did not include another cache line buffer in the MPC105 bridge simply because it believed that the performance of the MPC105 was sufficient, and adding another buffer (which would allow a transfer at a constant rate in the MPC105 system) would have been obvious.

11. The MPC105, as described by the evidence, was faster than the Viper described in the OPTi '906 patent.

404. During the MPC105 depositions, OPTi counsel asked whether the MPC105 inventors would have wanted the MPC105 to transfer data faster in MRM transactions. According to one of the Motorola papers, the MPC105 could transfer data in an MRM in the following cycles:

12. Summary of opinions to be expressed with respect to the Motorola MPC105.

407. The testimony of Michael Garcia, Brian Reynolds, and Christopher Bryant, is consistent with the text of the MPC105 User Manual and the numerous publications and presentations about the MPC105 that were presented to academic and commercial organizations in 1994. The testimony and documentary evidence shows that the design of the MPC105, including the implementation of the presnoop and prefetch features, was completed by at least March of 1994.

13. Claim chart & timeline.

408. A claim chart is attached hereto as Attachment H, illustrating that Motorola's prior invention renders the presnoop patents claims-in-suit invalid for anticipation or obviousness. A timeline is attached hereto as Attachment I.

D. The IBM prior art (MicroChannel TDB and Amini).

1. IBM Microchannel TDB (1993).

409. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are rendered obvious under § 103 by IBM's Micro Channel architecture described in the IBM TDB. As such, the asserted claims of the '906 Patent are invalid. The Micro Channel architecture includes a snoop of a second or next cache line during the transfer of the data in a first cache line, to reduce the latency caused during a burst transfer of data when the second cache line is not snooped until after the data in the first cache line is transferred. Thus, the architecture teaches predictive snooping.

410. IBM published a technical disclosure bulletin in October 1993, two years prior to OPTi's application for the '906 Patent. IBM's disclosure constitutes a 35 U.S.C. § 102(b) reference that cannot be sworn behind by OPTi. Since the disclosure renders obvious the subject matter of the '906 Patent, it renders the asserted claims of the patent invalid. A claim chart for the IBM Micro Channel art is attached hereto as Attachment J.

2. Amini (1993).

411. Claims 1, 7, 8, 9, 21, and 26 of the '906 Patent are anticipated under § 102 and/or rendered obvious under § 103 by the IBM Amini patent. U.S. Patent No. 5,581,714. (EX. 215).

412. IBM was one of the members of PCI SIG, and like fellow members Intel and Compaq, worked on a chipset to implement the PCI specification MRM command.

413. The IBM Amini patent describes a bridge chipset that includes a prefetch buffer that may be implemented in a "ping pong" structure so that data may be continuously transferred from the secondary memory to the bus master. Mr. Amini testified:

20

16 Q If you could look at Column 11 of
17 Exhibit 215. Beginning at line 6, it reads:
18 "The FIFO may be implemented,
19 for example, in a double word
20 (DWORD) interleaved structure or a 1:30:12PM
21 ping-pong structure."
22 Do you see that?
23 A Yes.
24 Q What is a ping-pong structure?
25 A I have -- I'd have to take a closer look 1:30:28PM

21

1 specifically to refresh my memory. 1:30:30PM
2 Q Okay.
3 A But more than likely, at a high level, it
4 probably references to being able to write into it
5 from one bus and read from it to the other bus at 1:30:49PM
6 the same time.
7 Q So would a ping-pong structure include two
8 structures, one that could write while the other was
9 reading at the same time?
10 MR. BRODY: Object to the form. 1:31:07PM
11 THE WITNESS: Can you ask your question
12 again?
13 BY MR. TETER:
14 Q Sure.
15 Would a ping-pong FIFO include two 1:31:13PM
16 components, one that could write while the other one
17 was reading?
18 MR. BRODY: Object to the form.
19 THE WITNESS: Probably, yes.

414. With a buffer implemented in a ping-pong structure, the Amini patent would allow bursting across cache line boundaries without delay, as the buffer would start to fill from the DRAM as soon as data units left the buffer for the PCI bus master and space became available. Mr. Amini testified:

66

24 What's your understanding, if any, of what
25 it means that data is transferred into the FIFO 2:47:51PM

67

1 immediately when space becomes available in the 2:47:54PM
2 FIFO?
3 MR. BRODY: Object to the form; asked and
4 answered.
5 THE WITNESS: As I read this, you know, my 2:48:01PM
6 sense would be that by that, it meant that that
7 location that's empty can be used to -- to transfer
8 additional data into.
9 And that would be consistent with a
10 concept of, you know, ping-pong or trying to improve 2:48:19PM
11 the efficiency of the other transfers between two
12 busses.

415. A claim chart for the Amini patent is attached hereto as Attachment K.

E. In addition to being anticipated, the claims in suit were obvious in view of the Intel (including Shanley), Compaq, Motorola and IBM prior art.

1. Differences between the invention and the prior art.

416. The obviousness instructions state that the differences between the prior art and the claimed invention should be considered as follows:

10.9.2 DIFFERENCES BETWEEN THE INVENTION OF THE CLAIMS AND THE PRIOR ART

In determining the differences between the invention covered by the patent claims and the prior art, you should not look at the individual differences in isolation. You must consider the claimed invention as a whole and determine whether or not it would have been obvious in light of all of the prior art.

In deciding whether to combine what is described in various items of prior art, you should keep in mind that there must be some motivation or suggestion for a skilled person to make the combination covered by the patent claims. You should also consider whether or not the prior art "teaches away" from the invention covered by the patent claims. The

question to be answered is: Would someone reading the prior art be discouraged from following the path taken by the inventor?

417. The prior art does not “teach away” from the OPTi presnoop invention.

In fact, the prior art, beginning with the Shanley text, clearly encouraged persons of skill in the art to practice snoop ahead.

2. Level of ordinary skill.

418. The instructions address the level of skill in the art as follows:

10.9.3 LEVEL OF ORDINARY SKILL

Obviousness is determined from the perspective of a person of ordinary skill in the art. This person is presumed to know all of the prior art, not just what the inventor may have known. When faced with a problem, this ordinary skilled person is able to apply his or her experience and ability to the problem and also to look to any available prior art to help solve the problem.

Factors to consider in determining the level of ordinary skill in the art include the educational level and experience of people working in the field, the types of problems faced by workers in the art and the solutions found to those problems, and the sophistication of the technology in the field.

419. I addressed the level of ordinary skill in the art earlier in my report. It is my opinion that the level of ordinary skill in the field was relatively high, as persons in the chipset design field typically had practical work experience that would have led them to practice the claimed presnoop invention.

3. The objective indications of obviousness support my opinion.

420. As noted above, I have been given the following instruction concerning the law of obviousness:

10.9.4 OBJECTIVE INDICATIONS CONCERNING OBVIOUSNESS

You also must consider what are referred to as objective indications of non-obviousness. Some of these indications of non-obviousness are:

1. Commercial success of products covered by the patent claims or made by a process covered by the patent claims.
2. A long-felt need for the invention.
3. Failed attempts by others to make the invention.
4. Copying of the invention by others in the field.
5. Unexpected results achieved by the invention.
6. Praise of the invention by the infringer or others in the field.
7. The taking of licenses under the patent by others.

8. Expressions of surprise by experts and those skilled in the art at the making of the invention.

9. The patentee proceeded contrary to the accepted wisdom of the prior art.

The presence of any of these objective indications may suggest that the invention was not obvious. These objective indications are only relevant to obviousness if there is a connection, or nexus, between them and the invention covered by the patent claims. For example, commercial success is relevant to obviousness only if the success of the product is related to a feature of the patent claims. If the commercial success is the result of something else, such as innovative marketing, and not to a patented feature, then you should not consider it to be an indication of non-obviousness.

421. Although OPTi has not identified any evidence that it plans to argue what constitutes “objective indications of obviousness” I will analyze the factors below.

- a. The OPTi Viper was not a commercial success—Intel’s Triton, which included both prefetching and presnooping, had far greater performance and was successful.**

422. OPTi’s Viper was not a commercial success. An August 1995 article in the Microprocessor Report article explained:
Opti Fights for Market Share

Opti's long-awaited Viper-M chip set has been in production since May 1995, about three months behind Triton and Apollo Plus. As Figure 3 shows, Viper-M's system architecture is not as straightforward as that of other chip sets: part of the ISA bus connects to the system controller, and the rest connects to the peripheral controller.

In addition to the common features for a Pentium-class chip set, Viper-m has a few unique features. It supports the Sony Sonic-2WP cache module and adaptive write-back, which means that if a CPU write hits in the secondary cache and is a page hit in main memory, the write will be treated as a write-through cycle. Viper-M can also run its PCI interface in either synchronous or asynchronous modes. In addition, it is the only chip set discussed in this article that supports a VL-bus slave device. The VL support, which complicates the Viper-M architecture, really belongs to 486-class systems. It seems unlikely that this feature will help Opti gain market share.

Opti has been very successful in the chip-set business for many years. Its total market share for desktop PCs was about 20% in 1994. Like VLSI Technology, Opti is losing market share this year for three reasons. First, the Viper-M product has been late to market. Second, the product does not bring enough differentiation in performance and function. In particular, its PCI performance is relatively poor: the sustained data throughout is

about 60 Mbytes/s, while all the other chip sets discussed in this article can sustain 100 Mbytes/s or more. Although Opti calls Viper-M a multimedia chip set, it does not have the features to separate it from others as a dedicated multimedia solution. Finally, Intel has simply taken away Opti's market share.

Ironically, Opti has always been thought of as a low-cost chip maker, but it cannot compete with companies like VIA on cost. Opti was once the number-one chip-set supplier to Taiwan, but Opti's presence in the Far East market, which consumes more than 50% of all chip sets worldwide, is not significant any more. Opti seems unable to sustain its profit margins by selling to that market. In the U.S., Intel has stolen major system OEMs like Packard Bell, Gateway, and Dell that used to be either Opti or VLSI customers.

[Yong Yao, Vendors fight for Pentium core-logic market; Intel is king of the land, most others are losing chip-set share (Cover story), MICROPROCESSOR REPORT (August 21, 1995) (NV 021342-021348)].

423. Triton's commercial success was due to many factors – not just presnooping. As the Viper demonstrated, presnooping alone did not ensure success. Intel exceeded Viper's performance by including prefetch buffers that OPTi did not have.

b. No long-felt need existed – the PCI specification was very new and bus masters did not initially use the MRM command.

424. The presnoop invention was conceived by Intel, Compaq, Motorola, and OPTi at a time when the PCI specification was very new. As a result, those in the industry believed that few bus masters would be trying to take advantage of the MRM command. For example, Mr. Reynolds testified, Motorola believed that there simply weren't bus master devices at the time that would use the MRM command.

144

3 Q. So was it -- was it your view at the time that
4 the speculative read feature allowed transactions to be
5 fast enough?

6 A. Yes. And at the time we also were not
7 convinced for the 105 that there were devices that were
8 using memory read multiple as a command. There
9 probably were, but it was not -- I don't think it was
10 prevalent. And so that also factored into our decision
11 to --

12 Q. Are you suggesting that at the time there
13 wasn't as much of a need for a bridge chip that would

14 allow for memory read multiple commands without any
15 kind of latency at all between data transfers?
16 A. Yes. Yes.


c. Others did not try and fail to make the invention.

425. Others in the field did not fail to make the invention. In fact, many of the competitors in the field implemented *both* presnooping and prefetching – it appears from OPTi’s patent and the articles cited above that OPTi’s Viper chipset implemented only presnooping, and as a result, had a wait state inserted between each data transfer and a very low data transfer rate.

d. No evidence suggests that anyone copied the Viper – OPTi trailed Intel and others in PCI chipset performance.

426. No evidence suggests that anyone copied OPTi. Rather, as described above in my report and in the cited article, OPTi was viewed as lagging in performance to Intel’s Triton.

427. Because Intel’s Triton and Compaq’s CMC/DMC sent data units on each PCI clock cycle, they could achieve a 133 MB/s data transfer rate even across cache line boundaries. (The data rate for the first cache line was less.) According to the MPC105 articles, the MPC105, which could send DWORDs at a 4-1-1-1-1-1-1-1 clock cycles, should have been able to achieve a data rate of approximately $133 \times (8/11)$ MB/s, or 96 MB/s. According to the OPTi ‘906 patent, the system would have a wait state between each data transfer. Thus, according to the ‘906 patent figures, the ‘906 patent should have been able to achieve a data rate of only $133/2$ or 67 MB/s. *See Figure, below.*



Performance Comparison for Steady-State MRM Transfer Across Cache Lines
In MB/s

e. The invention did not achieve unexpected results.

428. Presnooping does not achieve unexpected results – the results are entirely predictable and obvious.

f. The invention did not receive praise of others.

429. OPTi's Viper was not praised, but was criticized by commentators as late arriving to market and as having poor performance, much slower than Intel's Triton. The OPTi inventor, Mr. Ghosh, testified that he did not publish any awards or articles describing the presnoop feature.

170

22 Q. Did you ever publish any articles about the
23 pre-snoop feature that you describe and claim in the
24 '906 patent?

25 A. No.

171

1 Q. Did you ever win any awards for the
2 predictive snooping invention?

3 A. Awards from industry, you mean?

4 Q. Industry. Sure.

5 A. No.

6 Q. Did OPTi give you a bonus for the
7 predictive snooping method?

8 A. Yes.

9 Q. What was the bonus?
10 A. Long back. Maybe 2,000, maybe \$5,000.

430. Shanley and others did praise Intel's Triton – but the Triton succeeded and was praised for many reasons, not just presnoop. Triton achieved the performance gain by both presnooping and prefetching. The Viper, as described in the '906 patent, only presnooped, and its performance lagged far behind. Similarly, the Compaq CMC/DMC chipset also prefetched and presnooped, and achieved far greater performance than the '906 patent.

431. Intel and Compaq did file patent applications for the Intel streaming invention (combination of presnoop and prefetch elements) and the Compaq concurrency streaming invention (combination of presnoop, prefetch, and concurrency in other operations on other buses in the system), but those patent applications covered much more than the mere presnoop that OPTi has claimed in the '906 patent. Intel and Compaq each described and claimed a more complicated combination of elements that allowed Intel and Compaq to achieve the significant performance gains that they reported and attained. As a result, the fact that Intel and Compaq filed patent applications for their inventions does not show that OPTi's '906 presnoop claims were novel.

g. The taking of licenses under the patent by others.

432. I understand that AMD did not take a license from OPTi. VIA did not take a license. Apple did not take a license. I am informed that Compaq indicated, through Mr. Dukker, that it did not want a license. The only party licensed to the OPTi portfolio is Intel. I addressed the Intel agreement above, and have explained why it does not suggest that OPTi's presnoop patents are valid.

433. OPTi's decision not to sue VIA, which failed to take a license, is noteworthy. VIA indicated that its own prior art chipset, the VT82C505, practiced presnooping and would invalidate the OPTi presnoop patents. The parties exchanged timing diagrams and disputed whether the VIA chipset operated as VIA stated it did.

434. OPTi states that it found a VIA board and could not reproduce the tests that VIA performed. According to Mr. Marren, VIA was lying.

5 Mr. Ferrell indicated that he believed the
6 wave forms would conclusively proved that Via was
7 published and was practicing the '906 invention more
8 that one year prior to '906 filing by OPTi. What
9 did OPTi conclude after seeing the wave forms in
10 Exhibit 260?

11 MR. GRACEY: You can answer that to the
12 extent that was communicated to Via.

13 THE WITNESS: Yes. But we communicated to
14 Via, it was total unadulterated bogus.

236

1 Q So was it your belief at that time that
2 Via was lying to OPTi?

3 A Yes.

236

435. Mr. Marren testified that OPTi does not agree that the VIA prior art practiced presnoop, but his belief is based on a board that OPTi no longer has, and cannot find.

24 Q Has OPTi agreed now at this point that the
25 Via prior art chipset in fact did practice presnoop?

243

1 MR. GRACEY: Object to form.

2 THE WITNESS: The answer to that is
3 absolutely no.

4 BY MR. TETER:

5 Q And that's based on OPTi's own testing of
6 the board that you described earlier?

7 MR. GRACEY: Object to form.

8 THE WITNESS: Yes.

9 Q

10 BY MR. TETER:

11 Q And that board no longer exists at OPTi,
12 correct?

13 A We are not sure yet.

14 Q You're not sure?

15 A We are still looking.

16 Q The board may be at OPTi?

17 A No, it isn't at OPTi, but maybe somewhere
18 we can locate it. We are still looking. We are
19 doing a search to see if we can find that board or a
20 duplicate of it.

21 Q Where are you looking for it?

22 A In Canada and in the U.S.

23 Q Did you ask Via to return it?
 24 A I believe so.
 25 Q What was Via's response?
 244 1 A They would try.
 2 Q They would try?
 3 A They would try.
 245

**h. No experts expressed surprise at the making of the invention—
 presnoop was fully expected and inevitable.**

436. I am not aware of any expression of surprise by anyone of skill in the art at the making of OPTi's Viper. As noted above, the Viper was criticized by commentators in the field.

437. The OPTi inventor, Mr. Ghosh, testified that nobody told him they were surprised by the results.

171

13 Q. Did anyone ever tell you that they didn't
 14 think that the pre-snooping system and method could
 15 be done?
 16 A. I don't get your question.
 17 Q. Did anyone ever say, Hey, it surprised me
 18 that this worked. I didn't think it could be done?
 19 A. I have never heard such a comment.

**i. The patentee did not proceed contrary to the accepted wisdom
 of the prior art.**

438. The conventional wisdom, as set forth in the Shanley text, was that presnooping was a technique that could improve performance. The earliest chipsets, such as Saturn, Mercury, and Neptune, did not include snoop ahead for the reasons discussed at length above. However, the "conventional wisdom" held that read ahead, and snoop ahead, should be used if practical. [See Shanley Text (EX. 101); Design Guide (EX. 97) (advocating read ahead); PCI Spec. 2.0 (EX. 64)].

F. OPTi's embodiment of the invention was a generation behind Intel and Compaq.

439. Consistent with the testimony of former OPTi founder Jaswa, I have noted above that OPTi appears to have been behind in PCI chipset development.

440. Unlike both the Intel Triton and Compaq Triflex/PCI chipsets, the OPTi Viper-M Chipset does not appear to utilize the peak 132 MB/s throughput available to PCI bus masters. The PCI local bus runs at 33 MHz and utilizes a 32-bit (or 4-byte) data path. As a result, the PCI local bus is capable of transferring 132 million bytes per second (33,000,000 cycles per second at 4 bytes per cycle) when bursting data. [See also EX. 64, PCI Local Bus Specification, Revision 2.0].

1.5. PCI Local Bus Features and Benefits

The PCI Local Bus was specified to establish a high performance local bus standard for several generations of product. The PCI specification provides a selection of features that can achieve multiple price-performance points and can enable functions that allow differentiation at the system and component level. Features are categorized by benefit below:


High Performance • Transparent upgrade from 32-bit data path (132 MB/s peak) to 64-bit data path (264 MB/s peak).

[EX. 64 at OPTINVIDIA 001530 (PCI Local Bus Specification, Revision 2.0 at 5)].

441. The evidence cited herein establishes that both the Intel Triton and the Compaq Triflex/PCI chipsets were capable of transferring data in every PCI clock period, continuously and across cache line boundaries. The OPTi Viper-M chipset, however, does not appear to have been capable of transferring data in every PCI clock period.

442. For example, OPTi's '906 Patent shows timing diagrams with intermittent wait states between data transfers of the PCI bus. Figure 4 of OPTi's '906 patent purports to show "the operation of the system of FIG. 1 in a situation where a PCI master has requested a burst read access to an address at the beginning of a cache line-sized block in the secondary memory address space" [EX. 17, '906 Patent Col. 12:21-25]. Figure 4 clearly

shows intermittent wait states between 32-bit data transfers on the PCI bus. These wait states are shown by the non-assertion (high) of the TRDY signal line in Figure 4. Figure 4 thus shows that each 32-bit data transfer on the PCI bus is interrupted by a wait state when TRDY is not asserted.



[EX. 17, '906 Patent, Figure 4].

443. Moreover, this interpretation of Figure 4 is confirmed by the testimony of Subir Ghosh, one of the inventors listed on the '906 patent:

91

2 Q. If you could look at Figure 4 again in the
 3 patent -- the '906 patent. You don't have to look
 4 at the annotated version of it. You can just look
 5 at the patent itself.

6 A. Okay.

7 Q. In Figure 4, signal line TRDY indicates
 8 that one data unit is transferred, then another data
 9 unit is transferred, and then another data unit is
 10 transferred; is that correct?

11 A. This one indicates the data units being
 12 transferred are always happening with one wait state
 13 in between.


14 Q. The wait state is TRDY nonasserted;
 15 correct?

16 A. That is correct.

444. Documents produced by OPTi are consistent with the '906 patent's description of the OPTi system and Mr. Ghosh's testimony. For example, OPTi produced the

Viper-M 82C556M/82C557M/82C558M Data Book which describes the operation and features of the OPTi Viper-M chipset. [OPTINVIDIA 007132-007389].

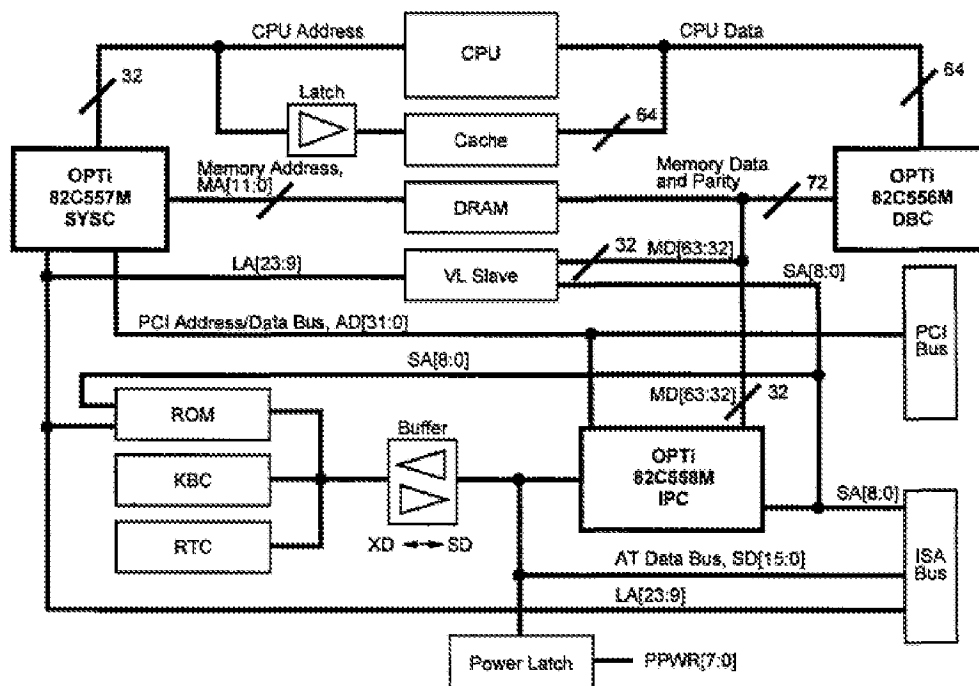
445. OPTi's interrogatory responses state that the Viper-M chipset practiced the presnoop claims. The Viper-M Data Book contains Figure 4-26, which purports to illustrate a "PCI Master Read from Local Memory":



446. Figure 4-26 also clearly shows that each 32-bit data transfer on the PCI bus is interrupted by a wait state.

447. The Viper-M Data Book (OPTINVIDIA 007132-007389) suggests why OPTi was unable to achieve the peak data transfer rate available on the PCI bus. In the Viper-M chipset, the PCI local bus was connected to the OPTi 82C558M IPC (Integrated Peripherals Controller) and received all data transfers from main memory (DRAM) via the 82C558M IPC. The system layout is illustrated in Figure 1-1 of the Viper-M Data Book.

Figure 1-1 System Block Diagram



[OPTINVIDIA 007154 (Viper-M Data Book at 1, Figure 1-1)].

448. As Figure 1-1 illustrates, the “PCI Bus” is connected to the “OPTi 82C558M IPC”. But as Figure 1-1 also illustrates, the 82C558M IPC is only connected to the upper 32 bits of the MD bus (MD[63:32]). Thus, for a low order DRAM read, the 82C558M needed the assistance of the 82C556M to capture the lower 32-bits of the DRAM read, and drive it back onto the upper half of the MD bus where the 82C558M could latch it and forward it onto the PCI bus:

For a low order DRAM read, the DRAM puts out the data on the MD[31:0]# lines. The 82C556M latches the data, inverts it and puts it out onto the MD[63:32] lines and drives it out to the 82C558M. The 82C558M then latches the data and puts it out on the AD[31:0] lines for the PCI master. For a low

[OPTINVIDIA 007240 (Viper-M Data Book at 87)].

449. The 82C556M could redirect the low order DRAM read because it contained a “memory to local bus buffer”:

2.1 82C556M (DBC) Data Buffer Controller

The 82C556M DBC performs the task of buffering the CPU to the DRAM memory data path. It also performs parity checking.

- CPU to memory data buffer
- CPU to local bus buffer
- Memory to local bus buffer

[OPTINVIDIA 007156 (Viper-M Data Book at 3)].

450. The 82C558M could handle high order DRAM transfers by itself because it was connected directly to the high order DRAM signal lines:

For a high order DRAM read, the DRAM puts out the data on the MD[63:32] lines. In this case, there is a direct path from the DRAM to the 82C558M and the 82C556M does not have to perform any latching or steering of data. The 82C558M latches the data available on the MD[63:32] lines and puts it out on the AD[31:0] lines for the PCI master. For a high order

[OPTINVIDIA 007240 (Viper-M Data Book at 87)].

451. However, the 82C558M could only receive 32-bits of data from the DRAM at a time (via the MD[63:32] signal lines), and does not appear to have had a data buffer where it could store additional data. [See OPTINVIDIA 007156 (Viper-M Data Book at 3) (listing features of 82C558M without buffers)]. The 82C558M block diagram also shows only 32 MD signal lines coming into the 82C558M and further confirms this operation:

Opti's long-awaited Viper-M chip set has been in production since May 1995, about three months behind Triton and Apollo Plus. As Figure 3 shows, Viper-M's system architecture is not as straightforward as that of other chip sets: part of the ISA bus connects to the system controller, and the rest connects to the peripheral controller.

In addition to the common features for a Pentium-class chip set, Viper-m has a few unique features. It supports the Sony Sonic-2WP cache module and adaptive write-back, which means that if a CPU write hits in the secondary cache and is a page hit in main memory, the write will be treated as a write-through cycle. Viper-M can also run its PCI interface in either synchronous or asynchronous modes. In addition, it is the only chip set discussed in this article that supports a VL-bus slave device. The VL support, which complicates the Viper-M architecture, really belongs to 486-class systems. It seems unlikely that this feature will help Opti gain market share.

Opti has been very successful in the chip-set business for many years. Its total market share for desktop PCs was about 20% in 1994. Like VLSI Technology, Opti is losing market share this year for three reasons. First, the Viper-M product has been late to market. Second, the product does not bring enough differentiation in performance and function. In particular, its PCI performance is relatively poor: the sustained data throughput is about 60 Mbytes/s, while all the other chip sets discussed in this article can sustain 100 Mbytes/s or more. Although Opti calls Viper-M a multimedia chip set, it does not have the features to separate it from others as a dedicated multimedia solution. Finally, Intel has simply taken away Opti's market share.

Ironically, Opti has always been thought of as a low-cost chip maker, but it cannot compete with companies like VIA on cost. Opti was once the number-one chip-set supplier to Taiwan, but Opti's presence in the Far East market, which consumes more than 50% of all chip sets worldwide, is not significant any more. Opti seems unable to sustain its profit margins by selling to that market. In the U.S., Intel has stolen major system OEMs like Packard Bell, Gateway, and Dell that used to be either Opti or VLSI customers.

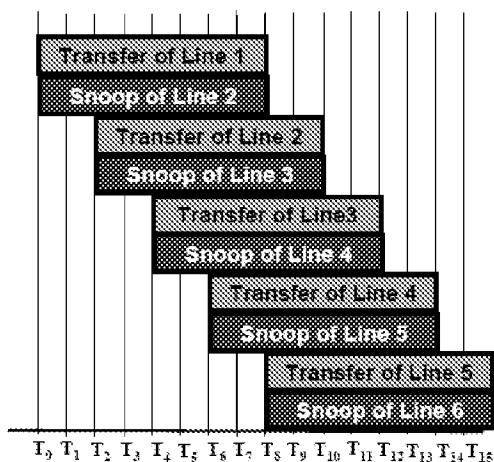
[Yong Yao, Vendors fight for Pentium core-logic market; Intel is king of the land, most others are losing chip-set share (Cover story), MICROPROCESSOR REPORT (August 21, 1995) (NV 021342-021348)].

454. The reason OPTi designed the Viper-M to have such relatively poor PCI performance, and the reason OPTi designed the '906 patent to have such poor performance, is not known, as I understand that OPTi did not retain all of its documents from the relevant time

frame. It is possible that because OPTi was a generation behind Intel and Compaq in PCI chipsets—as Mr. Jaswa testified—that OPTi designed the Viper in a hurried effort to have a PCI chipset on the market, even if it would be much slower than the chipsets that Intel and Compaq were designing. It is also possible that the OPTi designers ran out of gates on the die, or that OPTi had inexperienced designers who didn’t realize the impact of their design decisions. In any event, one thing appears clear from the evidence. During the 1994-95 time frame, Intel and Compaq’s chipsets were ahead of OPTi in terms of PCI performance.

G. Any claim that covers snooping ahead more than one line is not supported by a written description in the ‘906 patent specification.

455. OPTi has prepared the following Figure, which OPTi alleges shows how the claimed presnoop method and apparatus works. According to OPTi, the OPTi patent claims cover snooping ahead more than one line.



456. As set forth below, OPTi’s Figure is technically incorrect, and conflicts with the teaching of the patent.

1. The Transfer of Line N Cannot Occur Until the Snoop of Line N is Completed.

457. First, the transfer of an individual line cannot take place before the snoop of that line is complete. The whole point of snooping a line is to determine whether or not it is safe to even *begin* accessing the memory locations corresponding to that line – if there is a

snoop hit, that memory transfer cannot start until the dirty cache line has first been written to memory. OPTi's presnoop patents acknowledge that the snoop must be completed before the line transfers. [See, e.g., EX. 17, '906 Patent Col. 5:22-24 (noting that "every data transfer to or from the memory address space which is cached by the L1 cache should be preceded by an inquire cycle")]. Indeed, OPTi's claimed invention is to start the snoop early so that it will complete *before* the earliest possible start of transfer of the line to the PCI agent. Thus, the OPTi figure, which purports to show that the snoop of line 2 is occurring during the transfer of line 2, is incorrect and conflicts with the teaching of the presnoop patents. The snoop of a line quite simply cannot overlap with the transfer of the line, because the snoop must have completed before the transfer of that line can begin.

2. The Snoop of Line N+1 Cannot Occur Until the Snoop of Line N is Completed.

458. Second, OPTi's figure incorrectly shows that more than one line is being snooped at the same time, and between T₆ and T₇, purports to show *four* simultaneous snoops. Simultaneous snoops of different lines are not possible, because each snoop requires exclusive use of the host bus and exclusive use of the snoop port to the host's cache. Only one line may be snooped at a time, and nothing in the patent suggests snooping multiple lines at the same time. No CPU or cache had been invented at the time that allows for more than one line to be snooped at the same time, and nothing in OPTi's patent suggests one exists. Thus, contrary to OPTi's figure, the snoop of line N+1 cannot occur before the snoop of line N is completed.


3. The Transfer of Line N+1 Cannot Occur on the Bus Until the Transfer of Line N is Completed.

459. OPTi's Figure shows the transfer of multiple lines on the bus at the same time, where the green "transfer" boxes overlap in the horizontal (time) axis. In fact, only *part* of *one* line may be transferred on a bus at any particular time. More than one line of data cannot be transferred on the PCI bus at a given time – the PCI bus transmits an address for one

cache line, and the data for that line is then transferred to the bus master, one data unit at a time. It is impossible for different lines to be transferred on the PCI bus at a given time. Thus, contrary to OPTi's figure, the transfer of line N+1 on the bus cannot even begin before the transfer of line N is completed, and certainly cannot be concurrent with the transfer of line N.

4. OPTi's Patent Teaches Away from Snooping Ahead More than One Line.

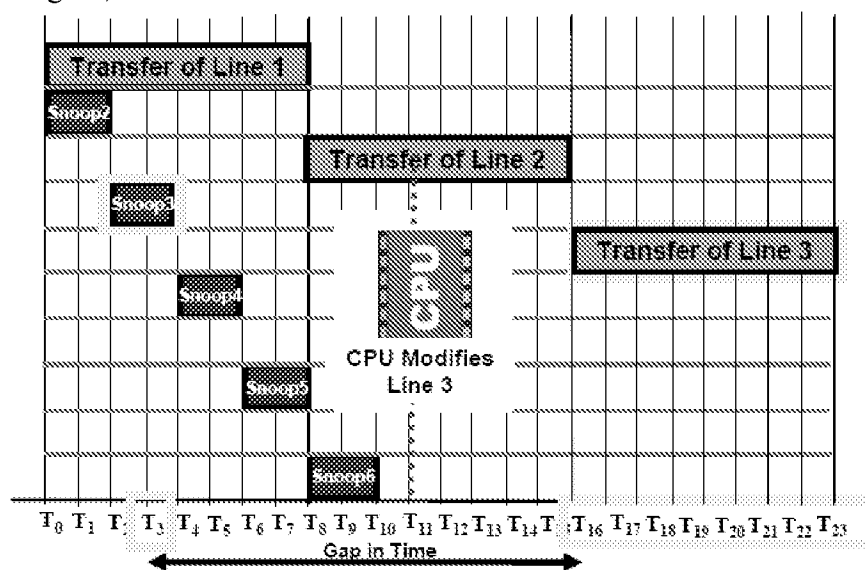
460. In essence, OPTi is attempting to expand its claims to cover snooping ahead more than one line ahead of the line being transferred on the PCI bus, in the manner shown below.



461. The dispute between the parties is whether OPTi's patent claims cover a system that snoops ahead more than the next line – *i.e.*, whether the patent would cover a system where the snoop of line N+2, N+3, or N+4 occurs during the transfer of line N on the PCI bus. In fact, the patent clearly teaches that only the *next line* (N+1) is snooped during the transfer of line N. Snooping ahead more than one line would defeat the purpose of the patent, and the patent teaches away from it. [See generally EX. 17, '906 Patent Col. 5:23-25 (noting that extra snoops “would severely hamper the performance of PCI masters performing burst cycles to or from secondary memory.”)]. Also note that the timing diagrams of '906 Figures 4-7 show only

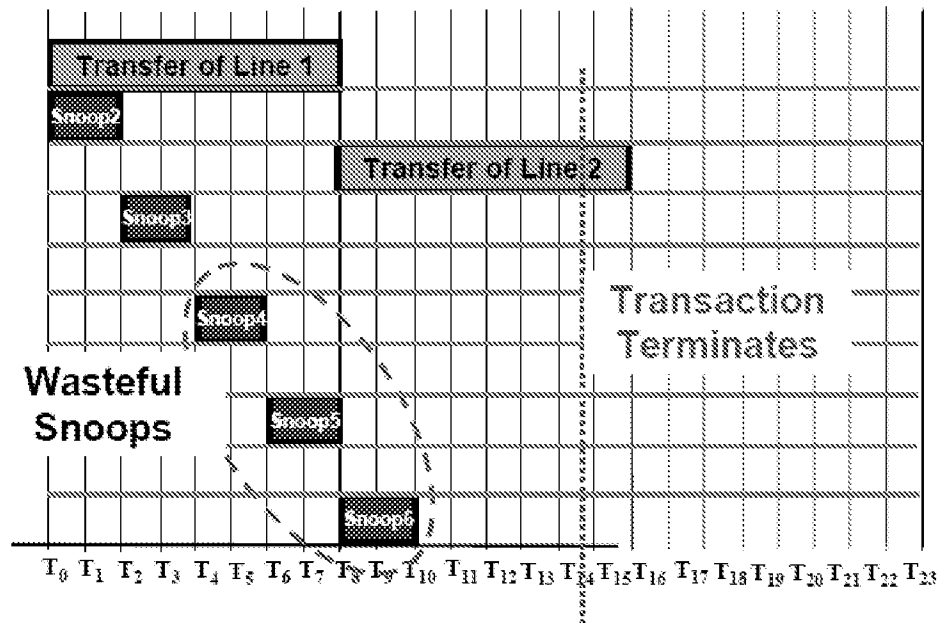
next-line presnooping, and the circuit diagrams are insufficient for any more aggressive forms of presnooping.

462. For example, if line 3 is being transferred many cycles after line 3 is snooped, as OPTi contends, the data in line 3 will be stale if the CPU modified the cache in the meantime. See Figure, below.




463. The OPTi patent clearly teaches that the “next-line” snoop ensures that the data will not be stale when it is transferred to the bus master. [See generally EX. 17, ‘906 Patent Col. 4:43-46 (“By this process, therefore, the external bus master will be able to access the desired line in main memory *without any further concern that the processor's L1 cache contains a more current copy of the data.*”)]. Thus, the OPTi patent clearly teaches away from snooping ahead more than one line, because snooping ahead more than one does not solve the problem, and does not ensure that the bus master receives valid data.

464. Similarly, the OPTi presnoop patents also teach that snoops take time and processor resources, and that unnecessary snoops should be avoided, which is something that persons of skill in the field at the time understood. [See generally EX. 17, ‘906 Patent Col. 5:23-25 (noting that extra snoops “would severely hamper the performance of PCI masters performing burst cycles to or from secondary memory.”)]. However, a system that snooped ahead more than one line would by definition perform many senseless snoops, as shown in the Figure below.



465. Where a snoop is initiated more than one line ahead of the line on the PCI bus, multiple senseless snoops would be performed, any or all of which could have instigated associated write backs. Such senseless snoops would burden the processor significantly, and the presnoop patent clearly teaches away from such a system. In fact, the OPTi patent discloses that in every embodiment, the snoop of line N+1 occurs while line N is being transferred on the PCI bus, as set forth at right. For the reasons set forth above, the patent clearly teaches away from any system that snoops more than one line ahead of the line being transferred on the PCI bus.



466. Thus, a system that snoops more than one line ahead of the line on the PCI bus is clearly not described by the OPTi presnoop patents. OPTi argues that persons of skill in the art would have understood that the presnoop invention could be implemented using buffers. Nothing in the patent suggests such a system. The OPTi patents teach away from and discourage any such system, because a system with buffers (1) would have to include another mechanism to make sure that the bus master receives valid data, and (2) would perform senseless snoops. Based on the clear and unambiguous discussion in the presnoop patents, persons of skill in the field would have understood that the presnoop patents teach away from snooping ahead more than one line. The patents describe speculatively snooping ahead only one line ahead of the line being transferred on the PCI bus.

467. I have been provided the following description of the law of written description:

A patent “must describe the invention sufficiently to convey to a person of skill in the art that the patentee had possession of the claimed invention at the time of the application, i.e., that the patentee invented what is claimed.” *LizardTech, Inc. v. Earth Resource Mapping, Inc.*, 424 F.3d 1336, 1345 (Fed. Cir. 2005).

To fulfill the written description requirement, the patent specification “must describe an invention and do so in sufficient detail that one skilled in the art can clearly conclude that the inventor invented the claimed invention.” *Lockwood v. American Airlines, Inc.*, 107 F.3d 1565, 1572 (Fed. Cir. 1997). “A patent applicant cannot disclose and claim an invention narrowly and then, in the course of an infringement suit, argue effectively

that the claims should be construed to cover that which is neither disclosed nor enabled in the patent.” *North American Vaccine, Inc. v. American Cyanamid Co.*, 7 F.3d 1571, 1577 (Fed. Cir. 1993). “It is not sufficient for purposes of the written description requirement of §112 that the disclosure, when combined with the knowledge in the art, would lead one to speculate as to modifications that the inventor might have envisioned, but failed to disclose.” *Lockwood v. American Airlines, Inc.*, 107 F.3d 1565, 1572 (Fed. Cir. 1997).

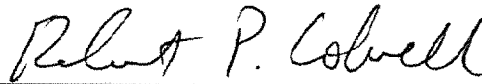
Recently, in *LizardTech*, the Federal Circuit affirmed a district court’s grant of summary judgment of invalidity for inadequate written description where a patent disclosed “only a single way” of performing the invention and “there is no evidence that the specification contemplates a more generic way.” *LizardTech*, 424 F.3d at 1344. The Federal Circuit noted that allowing a patentee to claim a broad, generic invention where the patent failed to describe it “would lead to sweeping, overbroad claims because it would entitle an inventor to a claim scope far greater than what a person of skill in the art would understand the inventor to possess or what a person of skill in the art would be enabled to make and use.” *LizardTech*, 424 F.3d at 134

468. The OPTi presnoop patent describe only a single way of predictive snooping – *i.e.*, one line ahead. For the reasons stated above, the patents teach away from anything else. Persons of skill in the art (Intel, Compaq, Motorola) also snooped only one line ahead, for the same reasons. It is my opinion that the OPTi patent does not provide a written description of any invention that snoops more than one line ahead.

XI. CONCLUSION AND RESERVATION.

469. For the reasons set forth above, in the attached claim charts, and in the attached timelines, it is my opinion that clear and convincing evidence shows that the asserted claims of the '906 patent are invalid due to both obviousness and anticipation. Although I have endeavored to provide a comprehensive discussion of my opinions and the basis for them, my analysis and study is ongoing, and I reserve the right to supplement this report or the opinions herein. In particular, I reserve the right to respond to any expert report provided by OPTi, and to address any new or additional information, evidence, or testimony that may be provided to me, including the testimony of fact or expert witnesses at trial and during expert depositions.

Dated: May 16, 2006



Robert P. Colwell, Ph.D